

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
25 May 2001 (25.05.2001)

PCT

(10) International Publication Number
WO 01/37473 A1

- (51) International Patent Classification⁷: **H04L 1/18**
- (21) International Application Number: **PCT/SE00/02268**
- (22) International Filing Date:
17 November 2000 (17.11.2000)
- (25) Filing Language: **English**
- (26) Publication Language: **English**
- (30) Priority Data:
09/441,882 17 November 1999 (17.11.1999) **US**
- (71) Applicant: **TELEFONAKTIEBOLAGET LM ERICSSON (publ)** [SE/SE]; S-126 25 Stockholm (SE).
- (72) Inventors: **JOHANSSON, Staffan**; Duvhöksgränd 15, S-974 55 Luleå (SE). **RIIHINEN, Wesa**; Kungsgatan 14, S-972 34 Luleå (SE).
- (74) Agent: **NORIN, Klas**; Ericsson Radio Systems AB, Common Patent Department, S-164 80 Stockholm (SE).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— *With international search report.*
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

(54) Title: **LINK LAYER ACKNOWLEDGEMENT AND RETRANSMISSION FOR CELLULAR TELECOMMUNICATIONS**

(57) Abstract: A mobile telecommunications system comprises a mobile user equipment unit (20) and at least one node (24) through which a packet switched data session is established between the user equipment unit and a data network (32). One of the node and the user equipment unit serves as a transmitter of segments of data packets and the other of the node and the user equipment unit serves as a receiver of the segments of data packets. If the first lost segment is not received within a timed interval, and if a second lost segment is also lost during the timed interval, at expiration of the timed interval the receiver sends a selective acknowledgment message (CSACK) to the transmitter which requests retransmission of both the first lost segment and the second lost segment. The acknowledgment messages of the system have a unique format including start sequence number and variable size bitmap field.



WO 01/37473 A1

LINK LAYER ACKNOWLEDGEMENT AND RETRANSMISSION FOR CELLULAR TELECOMMUNICATIONS

BACKGROUND

5

1. FIELD OF THE INVENTION

The present invention pertains to mobile telecommunications, and particularly to acknowledgment of receipt and retransmission of packet switched data for an upper layer protocol, such as the transmission control protocol/Internet protocol (TCP/IP), for
10 example.

2. RELATED ART AND OTHER CONSIDERATIONS

Initially commercial mobile or cellular telecommunications systems were primarily employed for voice calls, e.g., circuit switched connections. In more recent years, however, cellular telecommunications systems have also been employed for the
15 transmission of data (packet switched data), with the user equipment taking forms other than a mobile telephone. For example, user equipment such mobile laptops can send data over wireless links and through a cellular telecommunications system to wired computer networks such as the internet.

Cellular telecommunications systems employ a wireless link (e.g., air interface)
20 between the (mobile) user equipment unit and a base station (BS). The base station has transmitters and receivers for radio connections with numerous user equipment units. One or more base stations are connected to (e.g., by landlines or microwave) and managed by a radio network controller (RNC) [also known in some networks as a base station controller (BSC)]. The radio network controller is, in turn, connected through
25 control nodes to a core telecommunications network.

Control nodes can take various forms, depending on the types of services or networks to which the control nodes are connected. For connecting to connection-oriented, circuit switched networks such as PSTN and/or ISDN, the control node can be a mobile switching center (MSC). For connecting to packet switched data services
5 such as the Internet (for example), the control node can be a gateway data support node through which connection is made to the wired data networks, and perhaps one or more serving nodes. Examples of a particular packet data service called the General Packet Radio Service (GPRS) [provided in Europe in the context of the Global System for Mobile communications (GSM)] are provided by the following (all of which are
10 incorporated by reference): U.S. Patent Application Serial No. 09/069,969 filed April 30, 1998 entitled "Dynamic Allocation of Packet Data Channels"; U.S. Patent Application Serial No. Serial No. 09/069,939 filed April 30, 1998 entitled "Allocation of Channels for Packet Data Services"; and U.S. Patent Application Serial No. 09/090,186 filed June 4, 1998 entitled "Data Packet Radio Service With Enhanced
15 Mobility Management".

As indicated above, packet switched data services can include Internet service. In terms of Internet connection, the transmission control protocol/Internet protocol (TCP/IP) has gained wide acceptance. Although usually functioning together, the internet protocol (IP) and transmission control protocol (TCP) are actually separate
20 protocols, with the TCP being on a higher level (transport level) than the IP (on the network level).

There are numerous implementations of TCP, each with differing characteristics, the RENO implementation perhaps being the most common. In general, TCP supports a wide range of upper-layer protocols (ULPs). A ULP can send continuous streams of
25 data through TCP. The TCP breaks the streams into encapsulated segments, each segment including appropriate addressing and control information. TCP passes the segments to the network layer (e.g., the IP).

The IP layer encloses the TCP segments in IP packets or Internet datagrams. It is the Internet datagram that enables routing to source and destination TCPs in other
30 networks. Thus, the IP serves, e.g., to assemble IP datagrams and enable routing of the IP datagrams between IP addresses (e.g., between hosts) which are included in the IP datagram header.

TCP provides reliability which the IP lacks. In particular, the TCP carries out segmentation and reassembly functions of a datagram to match frame sizes and data-link layer protocols. In addition, TCP performs additional functions, such as addressing within a host, retransmission of lost packets, and flow control. General concepts
5 undergirding TCP/IP are understood from numerous publications, including Freeman, *Telecommunication System Engineering*, Third Edition, John Wiley & Sons, Inc., (1996), and W.R. Stevens, *TCP/IP Illustrated, Volume I: The Protocols* (Addison-Wesley, 1994).

Data losses because of bit errors occur over conventional wired links, but such
10 losses are so small as to be essentially non-existent (e.g., on the order of 10^{-6} over copper wire, and 10^{-9} over optical fiber). Such losses over conventional wired links stem almost exclusively from overflowing buffers in routers. TCP is designed to cope with these conditions, and consequently, packet losses are regarded as a congested network. Upon detection of loss, different implementations of TCP invoke different
15 congestion avoidance mechanisms, but generally all such congestion avoidance mechanisms decrease the transmission speed.

Some code-type error recovery capability (e.g., convolutional coding) is provided over the air interface, but such code-type error recovery cannot cope with large errors. Over the air interface, error recovery is performed locally with a local
20 retransmission protocol, wherein all data in a transmission buffer is cached until it has been successfully delivered to the receiver. In essence, any lost data is quickly transmitted by the local retransmission protocol before TCP has a chance to detect the loss. By retransmitting the data locally, faster recover can be done and, most importantly, the TCP will not detect the loss and accordingly will not invoke the TCP
25 congestion avoidance mechanism (unless data is lost somewhere other than over the air interface).

Thus, the task using the local retransmission protocol is how to realize quickly that data is lost, and how to retransmit the lost data. Traditionally, two primary types of strategies in local recovery have been utilized: (1) transport layer recovery of TCP
30 packets, and (2) link layer recovery of smaller data units called segments or frames (e.g., Automatic Repeat reQuest [ARQ]). Of the two general strategies, the ARQ-type of strategy is generally preferred.

Two examples of link layer recovery retransmission protocols are SSCOP (*see* ITU-T Recommendation Q.2110, 1994) and Radio Link Control (RLC). The SSCOP and RLC protocols are similar, with RLC being an improved variant of SSCOP. Fig. 13A illustrates a scenario of SSCOP operation, wherein segment S2_{13A} is lost. Receipt of the next segment S3_{13A} triggers a negative acknowledgment message (USTAT(S2)). Unfortunately, as illustrated in Fig. 13A, the negative acknowledgment message (USTAT(S2)) is lost. Moreover, segment S5_{13A} is also lost. Transmission of segment S6_{13A} therefore triggers a negative acknowledgment message (USTAT(S5)) to recover segment S5_{13A}. The negative acknowledgment message (USTAT(S5)) is successfully delivered to the transmitter, which then retransmits segment S5_{13A}. At this juncture, segment S2_{13A} is still not recovered. When a poll timer maintained by SSCOP expires, a poll message is transmitted to the receiver by setting a poll bit in a header of segment S9_{13A}. Upon reception of poll message S9_{13A}, the receiver checks its reception buffer, and notices that segment S2_{13A} is missing. Upon detection that segment S2_{13A} is missing, the transmitter transmits a STAT message, in particular STAT(S2). The STAT message is a selective acknowledgment message that can advise of gaps of one or more segments in the reception buffer. Upon reception of the STAT(S2) message, segment S2_{13A} is retransmitted.

Fig. 13B illustrates a similar scenario of RLC operation, wherein segment S2_{13B} is lost. Receipt of the next segment S3_{13B} triggers both transmission of a negative acknowledgment message (USTAT(S2)) and starting of an EPC timer (Estimated PDU Counter) to protect the USTAT message from loss. But the USTAT(S2) message from the receiver is lost, as is the segment S5_{13A} from the transmitter. The EPC timer expires in the third segment counting from where the USTAT(S2) message was sent, thereby triggering retransmission of the USTAT(S2) message. Although segment S5_{13B} has also been lost, recovery of segment S5_{13B} cannot begin until recovery of segment S2_{13B} is completed. Therefore, the EPC timer is stopped and restarted (for the recovery of segment S2_{13B}). When the USTAT(S2) message succeeds and segment S2_{13B} is received, the EPC timer is stopped and the reception buffer inspected. At this point segment S5_{13B} is noted as missing, and a recovery is initiated with the USTAT(S5) message. The recovery of segment S5_{13B} succeeds and the transmission of data packets can continue.

These two example link layer recovery retransmission protocols (SSCOP and RTL) can manage small losses fairly well, but as the Block Error Rate (BLER) grows, these two protocols are inadequate due to slow retransmission and transmission stalls resulting from filled transmission windows.

5 What is needed therefore, and an object of the present invention, is a link layer recovery technique which preserves good transmission utilization even during bad transmission conditions.

BRIEF SUMMARY OF THE INVENTION

10 A mobile telecommunications system comprises a mobile user equipment unit and at least one node through which a packet switched data session is established between the user equipment unit and a data network. One of the node and the user equipment unit serves as a transmitter of segments of data packets and the other of the node and the user equipment unit serves as a receiver of the segments of data packets. If a first lost segment is not retransmitted and successfully received within a timed
15 interval, and if a second lost segment is also lost during the timed interval, at expiration of the timed interval the receiver sends an acknowledgment message to the transmitter, requesting retransmission of both the first lost segment and the second lost segment.

 The acknowledgment messages of the system have a unique format which includes a starting sequence number field and a variable size bitmap. The start
20 sequence number (ssn) is mapped to a first non-received segment in a receive buffer; the first bit in the bitmap maps to a segment with a sequence number equal to the start sequence number. Bits in the bitmap are set to a first predetermined value to indicate receipt of respective ones of the plural segments or to a second predetermined value to indicate non-receipt of respective ones of the plural segments. In differing
25 embodiments, the bitmaps of the acknowledgment messages may have or may not have a stop bit. When a stop bit is used, the last bit set to the second predetermined value delineates the end of the information subset of the bitmap.

 The transmitter has a poll timer that is started/restarted and canceled in accordance with novel timer operational rules. In particular, starting/restarting a poll
30 timer in the transmitter upon occurrence of any of the following: (1) a sequence

number of a segment to be sent from the transmitter to the receiver is greater than or equal to $V(a) + \text{MaxWin} - \text{MaxP}$, wherein $V(a)$ is a sequence number of a first segment in a transmission window, MaxWin is the size of the transmission window, and MaxP is a predefined integer less than MaxWin ; (2) one or more segments are to be
5 retransmitted from the transmitter to the receiver if the poll timer is already running; and (3) a segment to be transmitted is the last segment in the buffer. The integer MaxP is dynamically determined, and can depend upon one of (1) a number of consecutive lost acknowledgment messages allowed; and (2) a frequency of negative acknowledgment messages. A time out value for the poll timer is also dynamically
10 determined.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, features, and advantages of the invention will be apparent from the following more particular description of preferred embodiments as illustrated in the accompanying drawings in which reference characters refer to the
15 same parts throughout the various views. The drawings are not necessarily to scale, emphasis instead being placed upon illustrating the principles of the invention.

Fig. 1 is a schematic view of an embodiment of a telecommunications system which utilizes the present invention.

Fig. 2 is a diagrammatic view of functionalities provided in a mobile user
20 equipment unit which participates in the telecommunications system of Fig. 1.

Fig. 3 is a schematic view of hardware components included in the mobile user equipment unit of Fig. 2.

Fig. 4 is a schematic view of a radio network controller (RNC) node in accordance with an embodiment of the invention.

25 Fig. 5 is a schematic view of a base station (BS) node in accordance with an embodiment of the invention.

Fig. 6 is a schematic view of a retransmission/acknowledgment unit according to an embodiment of the invention.

Fig. 7A is a state diagram showing operations on a transmitter side of the retransmission/acknowledgment unit of Fig. 6.

Fig. 7B is a state diagram showing operations on a receiver side of the retransmission/acknowledgment unit of Fig. 6.

5 Figs. 8A-8C are diagrammatic views depicting differing scenarios for three respective poll timer start/restart conditions.

Fig. 9A is a diagrammatic view of a first example format of a acknowledgment message CSACK according to the invention.

10 Fig. 9B is a diagrammatic view of a second example format of a acknowledgment message CSACK according to the invention.

Figs. 10A-10D are diagrammatic views depicting differing scenarios of segment transmission, acknowledgments, and segment retransmission according to modes of the invention.

15 Fig. 11 is a diagrammatic view showing steps involved in adaptive control of determining a parameter MaxP according to a mode of the invention.

Fig. 12A and Fig. 12B are diagrammatic views showing steps involved in adaptive control of determining time out values for a poll timer and a receiver timer according to a mode of the invention.

20 Fig. 13A and Fig. 13B are diagrammatic views depicting differing scenarios of segment transmission, acknowledgments, and segment retransmission according to prior art SSCOP and RLC protocols, respectively.

Fig. 14 is a matrix showing how MaxP can be set in accordance with the number of consecutive lost positive acknowledgment messages allowed

DETAILED DESCRIPTION OF THE DRAWINGS

25 In the following description, for purposes of explanation and not limitation,

specific details are set forth such as particular architectures, interfaces, techniques, etc. in order to provide a thorough understanding of the present invention. However, it will be apparent to those skilled in the art that the present invention may be practiced in other embodiments that depart from these specific details. In other instances, detailed descriptions of well known devices, circuits, and methods are omitted so as not to obscure the description of the present invention with unnecessary detail.

Fig. 1 shows a telecommunications network 18 in which a user equipment unit 20 communicates with one or more base stations 22 over air interface (e.g., radio interface) 23. Base stations 22 are connected by terrestrial lines (or microwave) to radio network controller (RNC) 24 [also known as a base station controller (BSC) in some networks]. The radio network controller (RNC) 24 is, in turn, connected through a control node known as the mobile switching center 26 to circuit-switched telephone networks (PSTN/ISDN) represented by cloud 28. In addition, radio network controller (RNC) 24 is connected to Serving GPRS Support Node (SGSN) 25 and through backbone network 27 to a Gateway GRPS support node (GGSN) 30, through which connection is made with packet-switched networks (e.g., the Internet, X.25 external networks) represented by cloud 32.

Gateway GRPS support node (GGSN) 30 provides the interface towards the external IP packet networks and X.25 networks. Gateway GRPS support node (GGSN) 30 translates data formats, signaling protocols and address information in order to permit communication between the different networks. Backbone network 27 is an Internet Protocol (IP) network. Serving GPRS Support Node (SGSN) 25 provides packet routing to an from a SGSN service area, and serves GPRS subscribers which are physically located within the SGSN service area. Serving GPRS Support Node (SGSN) 25 provides functions such as authentication, ciphering, mobility management, charging data, and logical link management toward the user equipment unit. A GPRS subscriber may be served by any SGSN in the network depending on location. The GPRS traffic is routed from the Serving GPRS Support Node (SGSN) 25 to base station controller (BSC) 24 and via base station (BS) 22 to user equipment unit 20. The functionality of Serving GPRS Support Node (SGSN) 25 and Gateway GRPS support node (GGSN) 30 may be combined in the same node, or may exist in separate nodes as shown in Fig. 1.

As understood by those skilled in the art, when user equipment unit 20 is participating in a mobile telephonic connection, signaling information and frames of user information from user equipment unit 20 are transmitted over air interface 23 on designated radio channels to one or more of the base stations 22. The base stations have
5 radio transceivers which transmit and receive radio signals involved in the connection or session. For information on the uplink from the user equipment unit 20 toward the other party involved in the connection, the base stations convert the radio-acquired information to digital signals which are forwarded to radio network controller (RNC)
10 24. The radio network controller (RNC) 24 orchestrates participation of the plural base stations 22 which may be involved in the connection or session, since user equipment unit 20 may be geographically moving and handover may be occurring relative to the base stations 22. On the uplink, radio network controller (RNC) 24 picks frames of user information from one or more base stations 22 to yield a connection between user
15 equipment unit 20 and the other party, whether that party be in PSTN/ISDN 28 or on the packet-switched networks (e.g., the Internet) 32.

One type of a user equipment unit 20 with which the present invention is particularly useful is a computer with mobile termination, such as a laptop computer, for example. An illustrative embodiment of a suitable user equipment unit 20 for the present invention is provided in Fig. 2. As shown in Fig. 2, user equipment unit 20 has
20 the following functional entities pertinent to the present invention: mobile termination entity (MT) 40; terminal adapter (TA) 42; terminal equipment 44; and a set 46 of applications. While each of these entities are described below, it should be understood that the invention is not confined to user equipment units having the same physical separation between functional entities, and that the present invention can be
25 implemented in other than the described functional configuration.

Mobile termination entity (MT) 40, which is sometimes called the Mobile Equipment (ME), contains the radio transmitter/receiver TX/RX 60 (with antenna 61) and communications control 62 toward the network 18, e.g., the setup and release of radio connections, handover, etc. Mobile termination entity (MT) 40 can be a standard
30 mobile pocket telephone (e.g., a GSM phone) or a phone card within user equipment unit 20.

Terminal adapter (TA) 42 acts as an adaptation between mobile termination entity (MT) 40 and the applications in the set 46 of applications. The terminal adapter (TA) 42 is typically realized as a Modem implemented on a PCMCIA (Personal Computer Memory Card International Association) card, which is inserted in a slot of terminal equipment 44. The terminal adapter (TA) 42 has a CPU 63 as well as a RAM 64 and a MT interface (I/F) 65.

Terminal equipment 44 is normally a small computer (or computer platform), and as such includes both hardware and software. Terminal equipment 44 thus has typical aspects of a computer platform, e.g., a processor an operating system and middleware (Internet protocol suits, for example), collectively illustrated by reference numeral 70 in Fig. 2. In addition, terminal equipment 44 has control logic 72 (executed by the processor) for controlling terminal adapter (TA) 42. Control logic 72 performs set-up and release of calls to and from the network 18.

As shown in Fig. 2, the set 46 of applications illustrated for the example embodiment includes an Internet browser 80; a file transfer program (FTP) 82; an E-mail program 84; and voice service 86.

Each application in set 46 is normally a program which is executed by the processor of terminal equipment 44 and which interacts with the user via, e.g., data input devices such as a keyboard and/or mouse and output or display devices. These applications typically can run on any personal computer (with or without radio access). The applications in set 46 use a number of application programming interfaces (APIs) towards the terminal equipment 44. One or several of these APIs is for communications with the network 18. Examples of APIs are Unix BSD Socket, WinSock or more telecom-specific APIs such as the Microsoft Intel Telephony API, AT&T, and Novell TSAPI or OnTheMove Mobile API. Thus, although the set 46 of applications is represented in Fig. 2 as an entity separate from terminal equipment 44, it should be understood that the set 46 of applications executed on the terminal equipment 44 if the terminal equipment 44 is a general computer, with the applications that are executed using the APIs offered by terminal equipment 44.

Fig. 3 shows how the functional entities of mobile station as illustrated in Fig. 2 are mapped onto hardware components of mobile station 20. In essence, Fig. 3 shows

terminal equipment 44 wherein mobile termination entity (MT) 40 and terminal adapter (TA) 42 are cards situated in card slots. Terminal adapter (TA) 42 is connected to central processing unit (CPU) 100 by bus 102. Mobile termination entity (MT) is connected to MT interface 65 of terminal adapter (TA) 42 by a cable.

5 Memories of terminal equipment 44, particularly read only memory (ROM) 104 and random access memory (RAM) 106 are also connected to central processing unit (CPU) 100 by bus 102. In RAM 106 are stored the TA control logic 72, the set 46 of applications, and TCP/IP stack 108.

10 Terminal equipment 44 interfaces with a user through input device(s) 110 and output device(s) 112, each connected through respective appropriate interfaces 120 and 122 to bus 102. Input device(s) 110 can be a keyboard and/or mouse, for example, while output device(s) 112 can take the form of a display device, such as a LCD display panel, for example.

15 Further details of an example radio network controller (RNC) 24 are shown in Fig. 4 as comprising a switch 240. Switch 240, which is controlled by RNC control unit 242, has a plurality of ports which are connected to various units and interfaces of radio network controller (RNC) 24. The RNC control unit 242 is connected to each element of radio network controller (RNC) 24, including timing unit 241; diversity handover unit (DHU) 245; and input/output unit 247. The radio network controller (RNC) 24 is
20 connected to a signaling network via a signaling interface 243. The interfaces connected to ports of switch 240 include service node interface unit 244; RNC I/F unit 246; and base station interface units 248. The service node interface unit 244 is connected to the appropriate mobile switching center 26 and GRPS control node 30. RNC interface unit 246 is connected by inter-RNC transport link 232 to other
25 (unillustrated) radio network controllers (RNCs) 24. Base station interface units 248 are connected to the base stations (BS) 22 served by radio network controller (RNC) 24. The input/output unit 247 and is connected to appropriate core ports of switch 240. The input/output unit 247 and a constituent channel switching unit thereof are described in more detail in companion United States Patent Application Serial No. 09/441,883
30 filed simultaneously herewith and entitled "CHANNEL SWITCHING IN MOBILE TELECOMMUNICATIONS WITH TCP CONSIDERATIONS", and which is incorporated herewith by reference.

Transmissions over air interface (air link) 23 are typically bidirectional. That is, some data packet sessions are transmitted from user equipment unit 20 over air interface 23 toward packet switched networks 32 (e.g., uplink relative to user equipment unit 20); other data packet sessions are transmitted in the opposite direction from packet switched networks 32 to user equipment unit 20 (downlink relative to user equipment unit 20). Accordingly, both user equipment unit 20 and radio network controller (RNC) 24 have a link layer retransmission/acknowledgment unit RAU. In this regard, and as mentioned above, user equipment unit 20 has retransmission/acknowledgment unit RAU_{UE} in its terminal adapter (TA) 42 (see Fig. 2 and Fig. 3). Similarly, the radio network controller (RNC) 24 has retransmission/acknowledgment unit RAU_{RNC} in packet I/O unit 247 (see Fig. 4).

Functionality included in a representative retransmission/acknowledgment unit RAU 600 is illustrated in Fig. 6. The retransmission/acknowledgment unit RAU 600 includes interface 602; link interface 604; transmission buffer 610; transmission controller 612; reception buffer 620; and reception controller 622. Depending on whether the retransmission/acknowledgment unit RAU 600 is employed in user equipment unit 20 or the node, the interface 602 is either an interface to the node or an interface to the remainder of the user equipment unit 20, respectively. Both transmission buffer 610 and reception buffer 620 are FIFO-type buffers. Segments being transmitted from interface 602 are stored in transmission buffer 610, under supervision of transmission controller 612, in route to link interface 604. Conversely, segments received from the link interface 604 are stored in reception buffer 620, under supervision of reception controller 622, in route to interface 602. Thus, at the moment shown in Fig. 6, transmission buffer 610 has TCP segments $S_{SSN(T)}$ through $S_{SSN+k(T)}$ stored therein, wherein in the subscript (1) SSN represents the start or first sequence number in the buffer, (2) "k" represents the size of the transmission buffer 610; and (3) the parenthetical T represents the transmit direction. The reception buffer 620 has TCP segments $S_{SSN(R)}$ through $S_{SSN+k(R)}$ stored therein, wherein in the subscript (1) SSN again represents the start or first sequence number in the buffer, (2) "k" represents the size of the reception buffer 620; and the parenthetical R represents the reception direction.

As illustrated in Fig. 6, transmission controller 612 includes, among other things, both a polling function 630 and a retransmission function 632 pertinent to the present invention. Basic operations performed by transmission controller 612, including polling

function 630 and retransmission function 632, are illustrated in Fig. 7A and described in conjunction therewith. The reception controller 622 includes an acknowledgment function 640 which is pertinent to the present invention. Similarly, basic operations performed by reception buffer 620 including the acknowledgment function 640 are
5 illustrated in Fig. 7B.

One purpose of acknowledgment function 640 is to generate acknowledgment messages for transmission to the transmitter. In the present invention, there are three basic types of acknowledgment messages, all of which are indicated by the general notation CSACK. As used herein, the notation CSACK may some times contain an
10 argument in the form CSACK(q), wherein q is a segment sequence number. The significance of the segment sequence number q depends on the type of acknowledgment message being sent. The three types of acknowledgment messages include: (1) a pure cumulative acknowledgment message in a normal situation (illustrated subsequently as message CSACK(a) in Fig. 8A); (2) a pure cumulative acknowledgment message sent
15 upon receipt of a last segment in a transmission buffer (illustrated subsequently as message CSACK(x) in Fig. 8C); and (3) a selective acknowledgment message which occurs upon detection of a lost segment. The first two types of acknowledgment messages CSACK are also collectively referred to as “positive” acknowledgment messages CSACK in that no segment losses have been detected. As used herein, the
20 phrase “transmission window” means a number of consecutive segments that can be transmitted before the transmitter requires receipt of an acknowledgment message from the receiver.

All three types of acknowledgment messages CSACK of the present invention preferably have the same format. An example of the format for one embodiment of
25 acknowledgment messages CSACK is illustrated in Fig. 9A. The acknowledgment message CSACK of Fig. 9A begins with a control message identification bit C in its first octet, which is followed in order by a start sequence number (ssn) and a bitmap. The start sequence number (ssn) bridges the first and second octets of the cumulative selective acknowledgment message. The bitmap is a variable size bitmap having bit
30 positions corresponding to segments. The first bit of the bitmap corresponds to the start sequence number (ssn), a last of the bits in the bitmap is mapped to a start sequence number in the start sequence number field. Thus, the following rules are followed for the format of the acknowledgment message CSACK Fig. 9A: (1) the start sequence

number (ssn) is always mapped to the first non-received segment in the receiver buffer; (2) a missing segment is denoted with a "1" in the bitmap; (3) received segments are denoted as "0" in the bitmap; and (4) the first bit in the bitmap maps to a segment with sequence number equal to the start sequence number (ssn). There is thus no stop bit in the bitmap of the acknowledgment message CSACK of Fig. 9A.

An example of the format for another embodiment of acknowledgment messages CSACK is illustrated in Fig. 9B. Like the format of Fig. 9A, the format of Fig. 9B has a control message identification bit C, followed by a start sequence number (ssn), followed by a bitmap. The four rules for the format of Fig. 9A also apply to the format of Fig. 9B. In addition, the format of Fig. 9B also has the following format rules: (5) the first "1" counting from the end of the bitmap denotes the end of the valid part of the bitmap (and thus serves as a stop sequence number field, e.g., a stop bit); and (6) if the acknowledgment message CSACK has a fixed size and the bitmap is too large to fit in the bitmap field, the bitmap is truncated and the last bit in the truncated bitmap is set to "1". Thus, in contrast to the format of Fig. 9A, the format of Fig. 9B does employ a stop bit.

The acknowledgment messages CSACK occurring in transmissions between the RAUs of the radio network controller (RNC) 24 and the user equipment unit 20 must consistent be one of format of Fig. 9A or of the format of Fig. 9B. That is to say, between two units acknowledgment messages CSACK cannot mix the formats of Fig. 9A and Fig. 9B. The stop bit of Fig. 9B is particularly beneficial in enabling the transmitter to release selectively buffered segments. This feature can be needed if buffer size is critical, for example in small mobile applications with stream type data transfer. The stop bit is not needed in transmissions of burst type, for example TCP.

To illustrate the benefit of a stop bit (of the format of Fig. 9B), consider the following scenario: the CSACK bitmap size is limited to eight bits (the first bit being the leftmost), and that segments 1, 2, 3, 4, and 6 - 15 have been successfully transmitted. An acknowledgment message CSACK without the stop bit (the format of Fig. 9A) would request retransmission of segment 5 with the following format: ssn=5, bitmap = 10000000. The only information this acknowledgment message CSACK can provide is that segment 5 has been lost, and that all segments up to and including segment 4 have been delivered. If the format of Fig. 9B (with the stop bit) is utilized,

on the other hand, the acknowledgment message CSACK is as follows: ssn = 5, bitmap = 10000001. The acknowledgment message CSACK of the format of Fig. 9B in this scenario not only tells that segment 5 is lost and that the segments up to and including segment 4 have been delivered, but also indicates that segments 6 - 11 are properly delivered and can be removed from the transmission buffer. Note that although segments 12 - 15 are properly delivered, this acknowledgment message CSACK cannot report that fact because of its short bitmap size [see rule (5)]. With a larger bitmap, however, segments 12 - 15 could have been acknowledged. For example, if the bitmap size for the format of Fig. 9B were instead sixteen bits, the content of the acknowledgment message CSACK would be as follows: ssn=5; bitmap = 10000000 00010000. This sixteen bit field enables the acknowledgment message CSACK to cover effectively all delivered segments and request retransmission of the lost segment 5.

In the ensuing discussion, the operations of Fig. 7A describe activities of transmission controller 612 for a transmitter retransmission/acknowledgment unit, while the operations of Fig. 7B pertain to activities of reception controller 622 for a receiver retransmission/acknowledgment unit. As mentioned above, depending upon the direction of the session, one of user equipment unit 20 and the radio network controller (RNC) 24 is the transmitter RAU, while the other of the user equipment unit 20 and the radio network controller (RNC) 24 is the receiver retransmission/acknowledgment unit RAU. Thus, although the following discussion of Fig. 7A and Fig. 7B refers to reference numerals of Fig. 6 with reference to both the transmitting and receiving operations, it should be understood that with respect to one direction of transmission the transmitting functions of Fig. 7A are performed in one equipment while the receiving functions of Fig. 7B are performed in another equipment.

Prior to discussing the operation of transmission controller 612 and reception controller 622, it should be mentioned that the transmission controller 612 in its polling function 630 has a poll timer. In general, the poll timer serves to ensure (1) that feedback is received from the receiver; and (2) that the last segment in a transmission window of the transmission buffer is delivered. The poll timer is never active unless a poll message has been sent. A poll message can be, and typically is, a segment in which a poll bit in the data header of the segment has been set. To avoid sending unnecessary poll messages, the poll timer must be canceled once the necessary information has

arrived. This means that the poll timer should not expire if the requested feedback was lost. In the present invention, the poll timer is started (or restarted) when any one of three start/restart conditions occurs. In addition, when the poll timer is started (or restarted) a poll bit is set in a data header of the segment. The poll timer is canceled or
5 stopped when either of two poll timer stop conditions occurs.

The first of the three alternative poll timer start/restart conditions is that the sequence number of the data packet to be transmitted is greater than or equal to Expression 1.

Expression 1: $V(a) + \text{MaxWin} - \text{MaxP}$.

10 In Expression 1, $V(a)$ is the sequence number of the first segment in the transmission window (e.g., the earliest or first un-acknowledged segment), MaxWin is the window size, and MaxP is a predefined integer less than MaxWin . In one test, MaxP was set to $\text{MaxWin}/2$, and MaxWin was set to 30. The first poll timer start/restart condition does not apply to restart the poll timer if it is already running.

15 A scenario depicting operation of the first poll timer start/restart condition is shown in Fig. 8A, wherein MaxWin is set to twelve and MaxP equals 5. The first non-acknowledgment segment is $V(a)$. As shown in Fig. 8A, when the transmission reaches segment $V(a) + 7$, the first poll timer start/restart condition is satisfied, so that the poll bit is set and the poll timer is started. When the acknowledgment message CSACK
20 returns, it is a cumulative acknowledgment with no loss indications and $\text{SSN} = V(a) + 7$.

The second poll timer start/restart condition is the retransmission of one or more segments if the poll timer is already running. A scenario depicting operation of the first poll timer start/restart condition is shown in Fig. 8B. In Fig. 8B, segment x is lost. The next received segment triggers a CSACK message, particularly message CSACK(x).

25 The EPC timer is then started. While the CSACK(x) message is in transit, the transmitter sends a poll message and starts a poll timer. Upon reception of the CSACK(x) message, the poll timer must be restarted so that the poll timer will not expire too soon. The poll bit of the retransmitted segment is also set. Note that a CSACK message is not sent on the poll message y , but instead segment x is awaited
30 before the message CSACK(z) is sent to acknowledge all received segments. When

CSACK(z) arrives, a second poll timer stop condition (discussed subsequently) is satisfied and the poll timer can be safely canceled.

The third poll timer start/restart condition is that the segment to be transmitted is the last segment in the buffer, e.g., LastTx = 1. A scenario depicting operation of the first poll timer start/restart condition is shown in Fig. 8C. In Fig. 8C, segment x is the last segment in the transmission buffer. Consequently, the LastTx flag in the segment is set. Accordingly, pursuant to the third poll timer start/restart condition, the poll timer bit is set and the poll timer is started. Upon reception of the acknowledgment message CSACK(x), a first poll time stop condition (discussed subsequently) is satisfied and the poll timer can be safely turned off.

The first poll timer stop condition occurs when (1) the last segment in the transmission buffer has been sent (e.g., LastTx=1), (2) a cumulative acknowledgment message has been received, and (3) the start sequence number field of the acknowledgment message CSACK is greater than the sequence number (Seqno) of the last transmitted segment. A scenario depicting operation of the first poll timer stop condition is shown in Fig. 8C, previously discussed.

The second poll timer stop condition occurs when (1) the segment to be transmitted is not a last segment in the transmission buffer (e.g., LastTx=0), and (2) any cumulative acknowledgment message is received which includes a start sequence number (ssn) which is greater than the first segment (V(a)) in the transmission window. Scenarios depicting operation of the second poll timer stop condition are shown in Fig. 8A and Fig. 8B.

The three alternative conditions which start or restart the poll timer, and the two stop conditions which cancel the poll timer, are herein also known as poll timer operational rules. The provision of the segmented acknowledgment message CSACK of the present invention, together with the poll timer and its associated poll timer operational rules, provides protection even with respect to pure cumulative acknowledgment messages and last segments transmitted from a transmission window.

Fig. 7A shows basic operations performed by and states assumed by transmission controller 612 in furtherance of the present invention. State 7A-1 depicts transmission

controller 612 as monitoring the filling of transmission buffer 610 with segments from interface 602, and then the reading out of segments from transmission buffer 610 to link interface 604. In similar manner, state 7B-1 shows reception controller 622 monitoring the filling of reception buffer 620 with segments from link interface 604, and the emptying of reception buffer 620 of segments being applied to interface 602.

Considering first the operation of the transmitter RAU as depicted in Fig. 7A, when either of the first or second poll timer start/restart rules (discussed above) are satisfied, the transmission controller 612 changes to state 7A-2 as indicated by event 7A-A in Fig. 7A. At state 7A-2 transmission controller 612 of the transmitter RAU sends a polling message to the receiver RAU. After sending the polling message of state 7A-2, transmission controller 612 automatically goes to state 7A-3 for the setting of the poll timer in its polling function 630, and then automatically returns to state 7A-1 for the monitoring of transmission buffer 610.

When the transmitter RAU receives an acknowledgment message CSACK from the receiver RAU as indicated by event 7A-D in Fig. 7A, transmission controller 612 goes from state 7A-1 to state 7A-4. At state 7A-4 transmission controller 612 processes the acknowledgment message CSACK. Upon receipt of the acknowledgment message CSACK, either of three courses of actions can be taken depending upon the type of acknowledgment message CSACK received (as indicated by its bitmap [see Fig. 9A or Fig. 9B]).

The first of the three potential course of action upon receipt of a pure cumulative acknowledgment message CSACK (positive CSACK) which occurs in a normal transmission situation. In other words, either of the two poll timer stop operation conditions/rules are met (as described above) and all segments up to the start sequence number specified in the acknowledgment message CSACK have already been released from transmission controller 612. In such case, as indicated by actions 7A-E transmission controller 612 stops the poll timer (at state 7A-5). The sending of the polling message is followed by restart of the poll timer at state 7A-3, and then transmission controller 612 resumes state 7A-1 as indicated by actions 7A-C and 7A-F.

A second of the three potential courses of action occurs upon receipt of a pure cumulative acknowledgment message CSACK triggered by receipt of the last segment

in a transmission window. This second course of action is implemented if the start sequence number specified in the acknowledgment message CSACK is greater than the sequence number of the last segment released from transmission controller 612. In this second course of action, as indicated by arrow 7A-G in Fig. 7A transmission controller 612 goes to state 7A-6 to release the already-delivered segments (up to the segment having the start sequence number identified in the ssn field of the acknowledgment message CSACK) from transmission buffer 610. If the bitmap of the acknowledgment message CSACK indicates no segment losses, transmission controller goes back to state 7A-1 as indicated by arrows 7A-J and 7A-F. If either of the poll timer stop rules is satisfied, transmission controller 612 stops the poll timer (at state 7A-5) before returning to state 7A-1.

If, upon release of segments from transmission buffer 610, the bitmap of the acknowledgment message CSACK indicates one or more segment losses, transmission controller 612 goes to state 7A-7 as indicated by arrow 7A-I. At state 7A-7 transmission controller 612 retransmits the lost segments from transmission buffer 610 to receiver RAU. From state 7A-7, transmission controller 612 either goes to state 7A-1 (as indicated by arrow 7A-L) when the poll timer second start/restart operation rule is not met, or to state 7A-2 (as indicated by arrow 7A-M) when the poll timer second start/restart operation rule is met.

A third of the three potential courses of action occurs upon receipt of a selective cumulative acknowledgment message CSACK (negative CSACK) having a bitmap which indicates loss of one or more segments. In such case, as indicated by arrow 7A-H, transmission controller 612 goes to state 7A-7 for retransmitting the lost segment(s) from transmission buffer 610 to receiver RAU. As indicated above, from state 7A-7, transmission controller 612 either goes to state 7A-1 (as indicated by arrow 7A-L) when the poll timer second start/restart operation rule is not met, or to state 7A-2 (as indicated by arrow 7A-M) when the poll timer second start/restart operation rule is met.

Considering next the operation of the receiver RAU as depicted in Fig. 7B, when monitoring the reception buffer 620 at step 7B-1, upon receipt of an in-order segment the reception controller 622 transitions to state 7B-2 as indicated by arrow 7B-B.

If it is noted, upon release at state 7B-2 of all in order-segments from reception

buffer 620, that there are no lost segments (or that all previously lost segments have been received), as indicated by arrow 7B-A reception controller 622 transitions to state 7B-4. At state 7B-4 reception controller 622 stops the receiver timer EPC, and then (as indicated by arrow 7B-C) returns to state 7B-1.

5 If, upon release at state 7B-2 of all in order-segments from reception buffer 620, it is noted that the last-received segment was a polling message (i.e., the poll bit is set in the received segment) and that the receiver timer EPC is not running, as indicated by arrow 7B-G the reception controller 622 transitions to state 7B-3. At state 7B-3, an acknowledgment message CSACK is generated by acknowledgment function 640 of
10 reception controller 622 and sent to transmitter RAU, after which reception controller 622 transitions back to state 7B-1 as shown by arrow 7B-H. In so doing, acknowledgment function 640 of reception controller 622 generates a positive acknowledgment message CSACK.

15 If, upon release at state 7B-2 of all in order-segments from reception buffer 620, it is noted that the last-received segment was a polling message (i.e., the poll bit is set in the received segment), that the receiver timer EPC is running, and that not all lost segments have been recovered, as indicated by arrow 7B-E the reception controller 622 transitions back to state 7B-1.

20 The reception controller 622 transitions directly from state 7B-1 to state 7B-3 (as indicated by arrow 7B-F) upon occurrence of any of the four following situations: (1) one or more lost segments have been detected and the receiver timer EPC is not running; or (2) the receiver timer EPC is running and something other than a retransmitted segment (RTX) arrives at the transmitter RAU; or (3) the EPC timer has expired; or (4) the EPC timer is running, one or more lost segments have been detected,
25 and a last retransmitted segment is received at transmitter RAU. The reception controller 622 transitions from state 7B-3 (the sending of the acknowledgment message CSACK) to state 7B-5 (as indicated by arrow 7B-I) when either (1) one or more lost segments have been detected and the receiver timer EPC is not running; or (2) the receiver timer EPC is running and something other than an RTX arrives at the
30 transmitter RAU. At state 7B-5, the receiver timer EPC is set. After setting of the receiver timer EPC, reception controller 622 transitions from state 7B-5

to state 7B-1 as indicated by arrow 7B-D.

Suppose that, in the process of the transmitter RAU sending segments from transmission buffer 610 to the receiver RAU, a segment is lost. For example, consider the scenario of Fig. 10A in which: (1) segment $S1_{10A}$ has been accurately transmitted from the transmitter to the receiver; (2) segment $S2_{10A}$ is lost (e.g., not received by the receiver); and (3) segment $S3_{10A}$ is received at the receiver RAU. In such a situation, in its monitoring of reception buffer 620 at state 7A-1 the reception controller 622 determines that a first lost segment is missing, i.e., segment $S2_{10A}$. Upon detecting the loss, reception controller 622 enters state 7B-3. In state 7B-3, the reception controller 622 of the receiver RAU sends a cumulative selective acknowledgment message $CSACK(S2)_{10A}$ to the transmitter RAU as indicated by message $CSACK(S2)_{10A}$ in Fig. 10A. In accordance with the format of Fig. 9B, the cumulative selective acknowledgment message $CSACK(S2)_{10A}$ of Fig. 10A would have the content: $ssn = (S2)_{10A}$; $bitmap = 101$. In the bitmap of the cumulative selective acknowledgment message $CSACK(S2)_{10A}$ the first "1" indicates loss of segment $(S2)_{10A}$, the "0" indicates that segment $(S3)_{10A}$ was successfully received; and the last "1" indicates the end of the bitmap. If no stop bit is used in accordance with the format of Fig. 9A, the bitmap will be 10000. Moreover, upon sending the cumulative selective acknowledgment message $CSACK(S2)_{10A}$, reception controller 622 enters state 7B-5 to set the receiver timer EPC prior to returning to state 7B-1.

Upon receipt of the cumulative selective acknowledgment message $CSACK(S2)_{10A}$ of Fig. 10A, the transmitter RAU determines that the last segment reported by the cumulative selective acknowledgment message $CSACK(S2)_{10A}$ is two segments after $S1_A$, i.e., $S3_A$. The sequence numbers for all segments corresponding to the bits of the bitmap can be ascertained relative to the start sequence number (SSN) carried in the cumulative selective acknowledgment message.

Receipt of the cumulative selective acknowledgment message $CSACK(S2)_{10A}$ of Fig. 10A at the transmitter RAU causes transmission controller 612 to enter state 7A-4 for processing of the cumulative selective acknowledgment message $CSACK(S2)_{10A}$ (see Fig. 7A). In processing the cumulative selective acknowledgment message $CSACK(S2)_{10A}$ at state 7A-4, reception controller 622 notes from the bitmap the loss of segment $S2_{10A}$, and accordingly transitions to state 7A-7 for retransmission of segment

S2_{10A}. After the retransmission of state 7A-7, the transmission controller 612 returns to state 7A-1. If the receiver RAU successfully receives the lost segment, reception controller 622 of the receiver RAU changes from state 7B-1 to state 7B-2, and upon noting release of all lost segments further stops the receiver timer EPC at state 7B-4 prior to returning to state 7B-1.

The situation described above with reference to Fig. 10A illustrates usage of the cumulative selective acknowledgment message CSACK to report loss of a single segment, e.g., usage of cumulative selective acknowledgment message CSACK(S2)_{10A} to report loss of segment S2_{10A}. The cumulative selective acknowledgment messages of the present invention can also report the loss of multiple segments, as illustrated in Fig. 10B. In particular, upon arrival of segment S5₁₀ in Fig. 10B, the reception controller 622 of the receiver RAU realizes that segments S2_{10B} through S4_{10B} have not been received into reception buffer 620. Therefore, as shown in Fig. 10B, a cumulative selective acknowledgment message CSACK(S2,S3,S4)_{10B} is sent at state 7B-3 from the receiver RAU to the transmitter RAU, with the cumulative selective acknowledgment message CSACK(S2,S3,S4)_{10B} including in its bitmap indications that each of segments S2_{10B} through S4_{10B} have not been received. In other words, the content of the cumulative selective acknowledgment message CSACK(S2,S3,S4)_{10B} for the format of Fig. 9B is ssn = (S2)_{10B}, bitmap = 1111000; and for the format of Fig. 9A is ssn = (S2)_{10B}, bitmap = 1110000. Again, the particular segments reference by each bit of the bitmap is determined with reference to the start sequence number (ssn) of the cumulative selective acknowledgment message CSACK(S2,S3,S4)_{10B}.

Fig. 10C depicts a situation which contrasts the operation of the present invention with the scenario depicted in Fig. 13A and Fig. 13B. In Fig. 10C, segment S2_{10C} is lost. Receipt of the next segment S3_{10C} triggers both transmission of a negative acknowledgment message (CSACK(S2)_{10C}) and starting of the receiver timer EPC to protect the (CSACK(S2)_{10C}) from lost. The negative acknowledgment message (CSACK(S2)_{10C}) is sent at state 7B-3, the timer EPC is set at state 7B-5 (see Fig. 7B). But the (CSACK(S2)_{10C}) message from the receiver RAU does get lost, as does the segment S5_{10C} from the transmitter. The EPC timer expires in the third segment counting from where the (CSACK(S2)_{10C}) message was sent. Upon detecting that the EPC timer has expired, reception controller 622 again enters state 7B-3 to generate another acknowledgment message CSACK, and resets the EPC timer (at state 7B-5).

The second acknowledgment message CSACK covers both lost segment $S2_{10C}$ and $S5_{10C}$, and is therefore denoted as $CSACK(S2, S5)_{10C}$. The acknowledgment message $CSACK(S2, S5)_{10C}$ has a content of $ssn = (S2)_{10C}$; $bitmap = 100101000$ for the format of Fig. 9B; and a $ssn = (S2)_{10C}$; $bitmap = 1001000$ for the format of Fig. 9A. Upon
 5 reception of the acknowledgment message $CSACK(S2, S5)_{10C}$, at state 7A-7 the transmission controller 612 retransmits the lost segments S2 and S5. Upon successfully receiving the retransmission of the segments S2 and S5, the reception controller 622 releases the segments S2 and S5 from reception buffer 620 and stops the receiver timer EPC at state 7B-4. Thereafter, the reception controller 622 returns to state 7B-1 for the
 10 processing of further segments.

The situation of Fig. 10C with its delaying of reporting of the loss of segment $S5_{10C}$ does result in delaying retransmission of segment $S5_{10C}$ for one or few frames. But as long as the transmission window is not closed, such delay does not affect overall performance. Depending on how many segments the preceding acknowledgment
 15 message concerns and what bandwidth is used, the delay can vary. Generally, the benefits of such bandwidth will more than offset the cost of such delay.

In the scenario of Fig. 10D, segment $S2_{10D}$ is lost, so that the reception of segment $S3_{10D}$ triggers a selective cumulative acknowledgment message $CSACK(S2)_{10D}$ [at state 7B-3] and starts receiver timer EPC [at state 7B-4]. Segments
 20 $S4_{10D}$ and $S5_{10D}$ are also lost, as is the acknowledgment message $CSACK(S2)_{10D}$. When the receiver timer EPC expires, a transition from state 7B-1 to state 7B-3 occurs. At state 7B-3, the reception controller 622 inspects reception buffer 620 and determines that segments $S2_{10D}$, $S4_{10D}$, and $S5_{10D}$ are all missing. At state 7B-3 reception controller 622 sends another selective cumulative acknowledgment message $CSACK(S2, S4, S5)_{10D}$
 25 and the receiver timer EPC is restarted. After retransmitted segments $S2_{10D}$, $S4_{10D}$, and $S5_{10D}$ are all received, the receiver timer EPC is canceled.

It should be noted that, in the Fig. 10D scenario, had $CSACK(S2)_{10D}$ been successfully transmitted and received, and the segment $S2_{10D}$ properly retransmitted and received, that the second selective cumulative acknowledgment message would have
 30 been message $CSACK(S4, S5)_{10D}$.

The foregoing illustrates a criteria of the present invention that, when a first

segment is lost and an acknowledgment message CSACK has been generated, no subsequent segment losses are recovered until (1) the first segment has been recovered, or (2) the receiver timer EPC has expired. If the first segment is recovered, the receiver timer EPC is canceled. If the receiver timer EPC has expired, it cannot be cancelled. In both cases, the reception buffer 620 is inspected, and if any segments (expressly including segments other than the first segment) are detected as lost, all lost segments are included in one acknowledgment message CSACK and recovery is re-initiated. The format of the acknowledgment message CSACK as explained in connection with Fig. 9A and Fig. 9B facilitates this operation and recovery.

The acknowledgment message format of Fig. 9A does not acknowledge gaps in the transmission window. In both the format of Fig. 9A and Fig. 9B, segments with sequence numbers up to and including the start sequence number (SSN) of the acknowledge message are released from the transmission buffer 310 upon reception of a cumulative selective acknowledgment message CSACK.

Various parameters of the invention can be made adaptable. As discussed below, these parameters include the MaxP value (employed in the Expression 1 of the first of the three alternative poll timer start/restart conditions), the time out or expiration value of the poll timer in transmission controller 612; and the time out or expiration value of the receiver timer EPC.

Recall that MaxP is a predefined integer less than MaxWin. It is advantageous to have an adaptive control scheme for determining the value of MaxP for several reasons. In a situation with high fault density on the air interface (e.g., many lost frames), setting MaxP to a high value means that there will be more time for recovery before the transmission stalls due to a full transmission window. On the other hand, with good radio conditions, setting MaxP to a low value means that feedback (e.g., acknowledgment messages) are not requested too often, thereby wasting less resources.

Thus, in one mode described with reference to Fig. 11, the present invention makes an estimate of the number of lost frames (the term "frame" being used interchangeably with "segment"), and based on that estimated value, sets the value of MaxP accordingly. As understood from the foregoing examples, a lost frame in the downlink direction (from the network to the user equipment) is detected by a gap in the

sequence numbers of the received segments. Using the acknowledgment message CSACK of the invention illustrated as 11-1 in Fig. 11, the user equipment requests the radio network controller (RNC) 24 to re-send the lost segments. The acknowledgment message CSACK includes, in its bitmap (see the format of Fig. 9A and Fig. 9B)

5 information regarding the number of segments that were lost since the last acknowledgment message CSACK was issued. The reception controller 622 can use this information regarding the number of lost segments to make an estimate of the block error rate (BLER). In this regard, and as shown as step 11-2 in Fig. 11, the estimate BLER(last) can be calculated as the number of lost segments as reported in the
10 acknowledgment message CSACK divided by the number of segments transmitted from the transmitter RAU to the receiver RAU since the previous acknowledgment message CSACK was received. Then, as reflected by step 11-3 in Fig. 11, in order to smooth out the estimate of the BLER(last), a standard type of low pass filter can be used, as understood with reference to Expression (2).

15 Expression (2): $BLER(est) = a * BLER(last) + (1-a) * BLER(old)$

In Expression (2), BLER(old) is next previous calculation of BLER(last), and a is a parameter in the range of from 0 to 1. The parameter a is employed to control the responsiveness of the BLER estimate. In this regard, setting a equal to zero means that BLER(est) is not updated at all; setting a equal to one means that BLER(est) is equal to
20 BLER(last). A value of a between zero and one gives a weighted sum of the last measured BLER [BLER(last)] and the previous estimate of BLER [e.g., BLER(old)]. Table 1 shows an example of calculating an on-going value of BLER(est) in which the parameter a is assigned a value 0.10 and the percentage of lost frames (i.e., BLER(last) alternates between 4% and 8%.

25 Using the BLER(est) value determined in a manner such as the foregoing, and as shown by step 11-3 in Fig. 11, a MaxP value can be selected. For example, transmission controller 612 can consult a lookup table and, using the BLER(est) value as an index, obtain a corresponding value for MaxP. The new MaxP value determined at step 11-4 can then be employed at step 11-5 as Expression 1 is evaluated to
30 determined when a new poll message should be sent (as step 11-6) to the user equipment (e.g., mobile station).

In general, the value for MaxP depends either on: (1) the number of consecutive lost positive acknowledgment messages allowed, keeping the risk for transmission stall below a certain limit (the lower the transmission stall limit is set, the larger MaxP grows); or (2) the frequency of negative acknowledgments (the higher the frequency of negative acknowledgments, the lower the MaxP value). Fig. 14 is a matrix showing how MaxP can be set in accordance with the number of consecutive lost positive acknowledgment messages allowed (i.e., criteria (1) mentioned above).

The risk for transmission stall (p) is calculated in accordance with Expression 3, in which N is set so that p is less than the limit. Expression (4) shows how to use N to calculate MaxP.

$$\text{Expression (3): } p = (\text{BLER})^N$$

$$\text{Expression (4): } \text{MaxP} = \text{BW/PS} * ((1+N) * \text{RTT} + N * \text{POLLT})$$

In Expression (4), BW is bandwidth (in bits/second); PS is segment size (in bits), RTT is round trip time (in seconds), and POLLT is poll timer value (in seconds). In the example using N values corresponding to the shaded field in the matrix of Fig. 14 are safe for a given BLER. In Fig. 14, RTT (round trip time) = 0.05 seconds; BW (bandwidth) = 32768 bits/second; PS (segment size) = 320 bits; and POLLT (poll timer value) = 0.04 seconds. For example, at 19% BLER, $N=3$ produces a safe MaxP value of 33. If the negative acknowledgments are taken into consideration, the MaxP value can be decreased.

The time out or expiration value of the poll timer in transmission controller 612 and the time out or expiration value of the receiver timer EPC can also be adaptively controlled, as illustrated in Fig. 12A. Thus far, no transport delays (except the delay introduced by the air interface) have been taken into consideration. However, delays can occur in the transport network as well as processing delays in the nodes involved (e.g., the base station node 22; the radio network controller (RNC) 24, and the user equipment). It is advantageous to set the time out values of the poll timer and receiver timer EPC to accommodate these transport network delays and processing delays.

As shown in Fig. 12A and Fig. 12B, the receiver RAU (e.g, the receiver in the

situation depicted in Fig. 12) measures the time interval from sending of a acknowledgment message CSACK (at step 12-1) to the time that a re-transmitted segment covered by acknowledgment message CSACK is received (step 12-4). The measured time interval is used as a value for TACK(last) [step 12-5], which is employed (at step 12-6) to develop an estimate TACK(est) using similar low pass filtering techniques as above described with reference to BLER. The estimate TACK(est) is then used for the new time out or expiration value for the receiver timer EPC.

The acknowledging strategy of the invention requires only one timer (EPC) in the receiver, but such timer should be adjustable accurately regardless of bandwidth changes and regardless of how many segments are to be retransmitted. Any type of timer works so long as it satisfies the criteria of adjustable accuracy regardless of (1) bandwidth changes, and (2) how many segments are to be retransmitted.

Concerning the time out value for the poll timer, the transmitter RAU (which, in the case of Fig. 12A, resides in the network) measures the time interval between the sending of a poll message and receipt of a corresponding acknowledgment message CSACK (see Fig. 12B). The measured time interval is used as a value TPOLL(last) [step 12-2], which is employed (at step 12-3) to develop an estimate TPOLL(est) also using similar low pass filtering techniques as above described with reference to BLER. The estimate TPOLL(est) is then used for the new time out or expiration value for the poll timer. It is to be noted that the measurements are re-initiated if the poll timer is restarted.

The present invention uses a polling acknowledgment strategy together with a loss-driven acknowledgment strategy. In accordance with the present invention, when segments are not lost very frequently, acknowledgment messages are sent often enough to provide good feedback, but yet still seldom enough to avoid unnecessary load on the acknowledgment link. Then, as the loss grows, loss driven acknowledgment takes over the responsibility of providing enough feedback to the transmitter. In this regard, the cumulative selective acknowledgment message provides feedback to the transmitter regarding which segments have been lost and which segment in the transmission buffer can be removed (e.g., released). To minimize band width needs, the cumulative selective acknowledgment message should generally be as small as possible.

The loss driven acknowledgment of the present invention thus provides bandwidth advantage. With an effective restraining mechanism, the protocol can keep down the amount of transmitted acknowledgment messages while still providing enough feedback. Combining the restraining mechanism with cumulative selective acknowledgment messages provides a robust and economic acknowledgment strategy.

By using the poll and acknowledgment strategy of the present invention, the performance of a link layer retransmission protocol can persist even during extremely lossy conditions. This means that satisfactory end-to-end TCP performance can be obtained even when the system is under very heavy load.

The aggressive acknowledgment cuts down the retransmission time of lost acks, especially in extreme-loss situations. The poll-ack mechanism ensures that no stalls except those resulting from massive loss (i.e., all segments are lost) will occur. Using a loss-adaptive polling mechanism minimizes the amount of requested acks at all times, while still holding the link utilization as high as possible.

The acknowledgment format of Fig. 9A makes buffer deallocation trivial as no selective algorithm is needed.

While the invention has been described in connection with what is presently considered to be the most practical and preferred embodiment, it is to be understood that the invention is not to be limited to the disclosed embodiment, but on the contrary, is intended to cover various modifications and equivalent arrangements included within the spirit and scope of the appended claims. For example, whereas the transmission controller 612 and reception controller 622 have herein been illustrated as two separate controllers for sake of clarity, it should be understood that the functions of transmission controller 612 and reception controller 622 can instead be performed by a single controller.

TABLE 1

BLER(last) %	BLER(old) %	BLER(est) %
0.00	0.00	0.00
4.00	0.00	0.40
8.00	0.40	1.16
4.00	1.16	1.44
8.00	1.44	2.10
4.00	2.10	2.29
8.00	2.29	2.86
4.00	2.86	2.97
8.00	2.97	3.48
4.00	3.48	3.53
8.00	3.53	3.98
4.00	3.98	3.98
8.00	3.98	4.38
4.00	4.38	4.34
8.00	4.34	4.71
4.00	4.71	4.64
8.00	4.64	4.97
4.00	4.97	4.88
8.00	4.88	5.19
4.00	5.19	5.07
8.00	5.07	5.36
4.00	5.36	5.23
8.00	5.23	5.50
4.00	5.50	5.35
8.00	5.35	5.62
4.00	5.62	5.46
8.00	5.46	5.71
4.00	5.71	5.54
8.00	5.54	5.79

WHAT IS CLAIMED IS:

1 1. A mobile telecommunications system comprising a mobile user equipment
2 unit (20); at least one node (24) through which a packet switched data session is
3 established between the user equipment unit and a data network (32), one of the node
4 and the user equipment unit serving as a transmitter of segments of data packets and the
5 other of the node and the user equipment unit serving as a receiver of the segments of
6 data packets;

7 characterized in that the receiver sends a first acknowledgment message
8 (CSACK) to the transmitter when the receiver detects loss of a first lost segment, and if
9 (1) the first lost segment is not received within a timed interval and (2) the receiver
10 determines that a second lost segment has been lost during the timed interval, at
11 expiration of the timed interval the receiver sends a second acknowledgment message to
12 the transmitter, the second selective acknowledgment message indicating loss of both
13 the first lost segment and the second lost segment.

1 2. The system of claim 1, wherein the first acknowledgment message is sent to
2 the transmitter upon receipt at the receiver of a segment which is sequenced subsequent
3 to the first lost segment.

1 3. The system of claim 1, further comprising a timer which times the timed
2 interval, the timer being adjustable.

1 4. The system of claim 1, wherein if the receiver determines that further lost
2 segments have been lost during the timed interval, the second selective
3 acknowledgment message also indicates loss of the further lost segments.

1 5. The system of claim 1, wherein at least one of the first acknowledgment
2 message and the second acknowledgment message includes a start sequence number
3 field and a bitmap field, wherein the start sequence number is mapped to a first non-
4 received segment; and wherein bits in the bitmap field are set to a first predetermined
5 value to indicate receipt of respective ones of the plural segments or to a second
6 predetermined value to indicate non-receipt of respective ones of the plural segments.

1 6. The system of claim 5, wherein the bitmap field of at least one of the first
2 acknowledgment message and the second acknowledgment message does includes a
3 stop sequence number field.

1 7. The system of claim 6, wherein the stop sequence number field is a stop bit.

1 8. The system of claim 6, wherein if the bitmap is too large to fit in the bitmap
2 field, the bitmap is truncated and the last bit in the truncated bitmap is set to the second
3 predetermined value.

1 9. The system of claim 5, wherein the transmitter has a poll timer, and wherein
2 the poll timer is started/restarted when any of the following occur:

- 3 (1) a sequence number of a segment to be sent from the transmitter to the
4 receiver is greater than or equal to $V(a) + \text{MaxWin} - \text{MaxP}$, wherein $V(a)$ is a
5 sequence number of a first segment in a transmission window, MaxWin is the
6 size of the transmission window, and MaxP is an integer less than MaxWin ;
7 (2) one or more segments are to be retransmitted from the transmitter to the
8 receiver if the poll timer is already running;
9 (3) a segment to be transmitted is the last segment in the buffer.

1 10. The system of claim 9, wherein the integer MaxP is dynamically
2 determined.

1 11. The system of claim 9, wherein the integer MaxP depends on a number of
2 consecutive lost acknowledgment messages allowed.

1 12. The system of claim 11, wherein MaxP is determined by the following
2 expression:

3
$$\text{MaxP} = \text{BW}/\text{PS} * ((1+N) * \text{RTT} + N * \text{POLLT})$$

4 wherein BW is bandwidth; PS is segment size, RTT is round trip time, and POLLT is a
5 poll timer value, and N is set to provide a risk for transmission stall p below a limit.

1 13. The system of claim 9, wherein the integer MaxP depends on a frequency of
2 negative acknowledgment messages.

1 14. The system of claim 9, wherein a time out value for the poll timer is
2 dynamically determined.

1 15. The system of claim 9, wherein the poll timer is canceled if:

- 2 (1) a segment to be transmitted is the last segment in the transmission buffer; and
3 (2) an acknowledgment message has been received which has a start sequence
4 number field which is greater than the sequence number of the last
5 transmitted segment.

1 16. The system of claim 9, wherein the poll timer is canceled if:

- 2 (1) a segment to be transmitted is not a last segment in the transmission buffer;
3 and
4 (2) any acknowledgment message is received which includes a start sequence
5 number which is greater than the first segment (V(a)) in the transmission
6 window.

1 17. A method of operating a mobile telecommunications system comprising at
2 least one node (24) through which a packet switched data session is established between
3 a user equipment unit (20) and a data network (32), the method including: using one of
4 the node and the user equipment unit to serve as a transmitter of segments of data
5 packets and using the other of the node and the user equipment unit to serve as a
6 receiver of the segments of data packets;

7 sending a first acknowledgment message from the receiver to the transmitter
8 when the receiver detects loss of a first lost segment;

9 characterized in that: when (1) the first lost segment is not received within a
10 timed interval and (2) a second lost segment is lost during the timed interval, the
11 method further comprises sending from the receiver, at expiration of the timed interval
12 sending a second selective acknowledgment message to the transmitter, the second
13 selective acknowledgment message indicating loss of both the first lost segment and the
14 second lost segment.

1 18. The method of claim 17, further comprising sending the first
2 acknowledgment message to the transmitter upon receipt at the receiver of a segment
3 which is sequenced subsequent to the first lost segment.

1 19. The method of claim 17, wherein a timer times the timed interval, and
2 wherein the method further comprises adjusting the timer.

1 20. The method of claim 17, further comprising, if further lost segments are lost
2 during the timed interval, using the second selective acknowledgment message also to
3 indicate loss of the further lost segments.

1 21. The method of claim 17, further comprising formatting at least one of the
2 first acknowledgment message and the second acknowledgment message to include a
3 start sequence number field and a bitmap field, wherein the start sequence number is
4 mapped to a first non-received segment; and wherein bits in the bitmap field are set to a
5 first predetermined value to indicate receipt of respective ones of the plural segments or
6 to a second predetermined value to indicate non-receipt of respective ones of the plural
7 segments.

1 22. The method of claim 21, further including formatting the bitmap field of at
2 least one of the first acknowledgment message and the second acknowledgment
3 message to include a stop sequence number field.

1 23. The method of claim 22, further including formatting the bitmap field to
2 include a stop bit as the stop sequence number field.

1 24. The method of claim 21, further including formatting the bitmap field
2 whereby if the bitmap is too large to fit in the bitmap field, the bitmap is truncated and
3 the last bit in the truncated bitmap is set to the second predetermined value.

1 25. The method of claim 17, wherein the transmitter has a poll timer, and
2 further comprising starting/restarting the poll timer upon occurrence of any of the
3 following:

- 4 (1) a sequence number of a segment to be sent from the transmitter to the
5 receiver is greater than or equal to $V(a) + \text{MaxWin} - \text{MaxP}$, wherein $V(a)$ is a
6 sequence number of a first segment in a transmission window, MaxWin is the
7 size of the transmission window, and MaxP is a predefined integer less than
8 MaxWin ;

- 9 (2) one or more segments are to be retransmitted from the transmitter to the
10 receiver if the poll timer is already running;
11 (3) a segment to be transmitted is the last segment in the buffer.

1 26. The method of claim 25, wherein the integer MaxP is dynamically
2 determined.

1 27. The method of claim 25, wherein the integer MaxP depends on a number of
2 consecutive lost acknowledgment messages allowed.

1 28. The method of claim 27, wherein MaxP is determined by the following
2 expression:

3
$$\text{MaxP} = \text{BW/PS} * ((1+N) * \text{RTT} + N * \text{POLLT})$$

4 wherein BW is bandwidth; PS is segment size, RTT is round trip time , and POLLT is a
5 poll timer value, and N is set to provide a risk for transmission stall p below a limit.

1 29. The method of claim 25, wherein the integer MaxP depends on a frequency
2 of negative acknowledgment messages.

1 30. The method of claim 25, wherein a time out value for the poll timer is
2 dynamically determined.

1 31. The method of claim 20, further comprising canceling the poll timer if:

- 2 (1) a segment to be transmitted is the last segment in the transmission buffer; and
3 (2) an acknowledgment message has been received which has a start sequence
4 number field which is greater than the sequence number of the last
5 transmitted segment.

1 32. The method of claim 20, further comprising canceling the poll timer if:

- 2 (1) a segment to be transmitted is not a last segment in the transmission buffer;
3 and
4 (2) any acknowledgment message is received which includes a start sequence
5 number which is greater than the first segment (V(a)) in the transmission
6 window.

1 33. A mobile telecommunications system which includes: a mobile user
2 equipment unit (20); at least one node (24) through which a packet switched data
3 session is established between the user equipment unit and a data network (32), one of
4 the node and the user equipment unit serving as a transmitter of segments of data
5 packets and the other of the node and the user equipment unit serving as a receiver of
6 the segments of data packets; and

7 characterized in that the receiver sends an acknowledgment message (CSACK)
8 to the transmitter upon occurrence of a predetermined condition, the acknowledgment
9 message including a start sequence number field and a bitmap field, the start sequence
10 number being mapped to a first non-received segment; and bits in the bitmap field being
11 set to a first predetermined value to indicate receipt of respective ones of the plural
12 segments or to a second predetermined value to indicate non-receipt of respective ones
13 of the plural segments.

1 34. The system of claim 33, wherein the bitmap field of at least one of the first
2 acknowledgment message and the second acknowledgment message does includes a
3 stop sequence number field.

1 35. The system of claim 34, wherein the stop sequence number field is a stop
2 bit.

1 36. The system of claim 33, wherein if the bitmap is too large to fit in the
2 bitmap field, the bitmap is truncated and the last bit in the truncated bitmap is set to the
3 second predetermined value.

1 37. The system of claim 33, wherein the predetermined condition is detection of
2 a lost segment.

1 38. The system of claim 33, wherein the predetermined condition is receipt at
2 the receiver of a poll message.

1 39. The system of claim 33, wherein the predetermined condition is a receipt of
2 a full transmission window of segments at the receiver.

1 40. A method of operating a mobile telecommunications system comprising at
2 least one node (24) through which a packet switched data session is established between
3 a user equipment unit (20) and a data network (32), the method including: using one of
4 the node and the user equipment unit to serve as a transmitter of segments of data
5 packets and using the other of the node and the user equipment unit to serve as a
6 receiver of the segments of data packets;

7 sending an acknowledgment message (CSACK) from the receiver to the
8 transmitter upon occurrence of a predetermined condition;

1 characterized by: formatting the acknowledgment message to include a starting
2 sequence number field and a bitmap field, the start sequence number being mapped to a
3 first non-received segment; and bits in the bitmap field being set to a first predetermined
4 value to indicate receipt of respective ones of the plural segments or to a second
5 predetermined value to indicate non-receipt of respective ones of the plural segments.

1 41. The method of claim 40, further including formatting the bitmap field of at
2 least one of the first acknowledgment message and the second acknowledgment
3 message to include a stop sequence number field.

1 42. The method of claim 41, further including formatting the bitmap field to
2 include a stop bit as the stop sequence number field.

1 43. The method of claim 40, further including formatting the bitmap field
2 whereby if the bitmap is too large to fit in the bitmap field, the bitmap is truncated and
3 the last bit in the truncated bitmap is set to the second predetermined value.

1 44. The method of claim 40, wherein the predetermined condition is detection
2 of a lost segment.

1 45. The method of claim 40, wherein the predetermined condition is receipt at
2 the receiver of a poll message.

1 46. The method of claim 40, wherein the predetermined condition is a receipt of
2 a full transmission window of segments at the receiver.

1 47. A mobile telecommunications system includes: a mobile user equipment
2 unit (20); at least one node (24) through which a packet switched data session is
3 established between the user equipment unit and a data network (32), one of the node
4 and the user equipment unit serving as a transmitter of segments of data packets and the
5 other of the node and the user equipment unit serves as a receiver of the segments of
6 data packets; and

7 characterized by: a poll timer and a buffer being included in the transmitter, and
8 the poll timer being started/restarted upon occurrence of any of the following occur:

9 (1) a sequence number of a segment to be sent from the transmitter to the
10 receiver is greater than or equal to $V(a) + \text{MaxWin} - \text{MaxP}$, wherein $V(a)$ is a
11 sequence number of a first segment in a transmission window, MaxWin is the
12 size of the transmission window, and MaxP is a predefined integer less than
13 MaxWin ;

14 (2) one or more segments are to be retransmitted from the transmitter to the
15 receiver if the poll timer is already running;

16 (3) a segment to be transmitted is the last segment in the buffer.

1 48. The system of claim 47, wherein the integer MaxP is dynamically
2 determined.

1 49. The system of claim 47, wherein the integer MaxP depends on a number of
2 consecutive lost acknowledgment messages allowed.

1 50. The system of claim 49, wherein MaxP is determined by the following
2 expression:

3
$$\text{MaxP} = \text{BW}/\text{PS} * ((1+\text{N}) * \text{RTT} + \text{N} * \text{POLLT})$$

4 wherein BW is bandwidth; PS is segment size, RTT is round trip time, and POLLT is a
5 poll timer value, and N is set to provide a risk for transmission stall p below a limit.

1 51. The system of claim 47, wherein the integer MaxP depends on a frequency
2 of negative acknowledgment messages.

1 52. The system of claim 47, wherein a time out value for the poll timer is
2 dynamically determined.

1 53. The system of claim 47, wherein the poll timer is canceled if:

- 2 (1) a segment to be transmitted is the last segment in the transmission buffer; and
3 (2) an acknowledgment message has been received from the receiver which has a
4 start sequence number field which is greater than the sequence number of the
5 last transmitted segment.

1 54. The system of claim 47, wherein the poll timer is canceled if:

- 2 (1) a segment to be transmitted is not a last segment in the transmission buffer;
3 and
4 (2) any acknowledgment message is received which includes a start sequence
5 number which is greater than the first segment ($V(a)$) in the transmission
6 window.

1 55. A method of operating a mobile telecommunications system comprising at
2 least one node (24) through which a packet switched data session is established between
3 a user equipment unit (20) and a data network (32), the method including: using one of
4 the node and the user equipment unit to serve as a transmitter of segments of data
5 packets and using the other of the node and the user equipment unit to serve as a
6 receiver of the segments of data packets, the transmitter having a buffer; and
7 characterized by: starting/restarting a poll timer in the transmitter upon
8 occurrence of any of the following:

- 9 (1) a sequence number of a segment to be sent from the transmitter to the
10 receiver is greater than or equal to $V(a) + \text{MaxWin} - \text{MaxP}$, wherein $V(a)$ is a
11 sequence number of a first segment in a transmission window, MaxWin is the
12 size of the transmission window, and MaxP is a predefined integer less than
13 MaxWin ;
14 (2) one or more segments are to be retransmitted from the transmitter to the
15 receiver if the poll timer is already running;
16 (3) a segment to be transmitted is the last segment in the buffer.

1 56. The method of claim 55, wherein the integer MaxP is dynamically
2 determined.

1 57. The method of claim 55, wherein the integer MaxP depends on a number of
2 consecutive lost acknowledgment messages allowed.

1 58. The method of claim 57, wherein MaxP is determined by the following
2 expression:

$$\text{MaxP} = \text{BW/PS} * ((1+\text{N}) * \text{RTT} + \text{N} * \text{POLLT})$$

3 wherein BW is bandwidth; PS is segment size, RTT is round trip time , and POLLT is a
4 poll timer value, and N is set to provide a risk for transmission stall p below a limit.
5

1 59. The method of claim 55, wherein the integer MaxP depends on a frequency
2 of negative acknowledgment messages.

1 60. The method of claim 55, wherein a time out value for the poll timer is
2 dynamically determined.

1 61. The method of claim 55, further comprising canceling the poll timer if:
2 (1) a segment to be transmitted is the last segment in the transmission buffer; and
3 (2) an acknowledgment message has been received which has a start sequence
4 number field which is greater than the sequence number of the last
5 transmitted segment.

1 62. The method of claim 55, further comprising canceling the poll timer if:
2 (1) a segment to be transmitted is not a last segment in the transmission buffer;
3 and
4 (2) any acknowledgment message is received which includes a start sequence
5 number which is greater than the first segment (V(a)) in the transmission
6 window.

1 63. A mobile telecommunications system including: a mobile user equipment
2 unit (20); at least one node (24) through which a packet switched data session is
3 established between the user equipment unit and a data network (32), one of the node
4 and the user equipment unit serving as a transmitter of segments of data packets and the
5 other of the node and the user equipment unit serving as a receiver of the segments of
6 data packets, the receiver sending acknowledgment messages;

7 characterized in that: the transmitter has a poll timer, and the poll timer being
8 started/restarted when any of the following occur:

9 (1) a sequence number of a segment to be sent from the transmitter to the
10 receiver is greater than or equal to $V(a) + \text{MaxWin} - \text{MaxP}$, wherein V(a) is a

- 11 sequence number of a first segment in a transmission window, MaxWin is the
12 size of the transmission window, and MaxP is an integer less than MaxWin;
13 (2) one or more segments are to be retransmitted from the transmitter to the
14 receiver if the poll timer is already running;
15 (3) a segment to be transmitted is the last segment in the buffer.

1 64. The system of claim 63, wherein the integer MaxP is dynamically
2 determined.

1 65. The system of claim 63, wherein the integer MaxP depends on a number of
2 consecutive lost acknowledgment messages allowed.

1 66. The system of claim 64, wherein MaxP is determined by the following
2 expression:

3
$$\text{MaxP} = \text{BW/PS} * ((1+N) * \text{RTT} + N * \text{POLLT})$$

4 wherein BW is bandwidth; PS is segment size, RTT is round trip time, and POLLT is a
5 poll timer value, and N is set to provide a risk for transmission stall p below a limit.

1 67. The system of claim 63, wherein the integer MaxP depends on a frequency
2 of negative acknowledgment messages.

1 68. The system of claim 63, wherein a time out value for the poll timer is
2 dynamically determined.

1 69. A method of operating a mobile telecommunications system comprising at
2 least one node (24) through which a packet switched data session is established between
3 a user equipment unit (20) and a data network (32), the method including: using one of
4 the node and the user equipment unit serving as a transmitter of segments of data
5 packets and using the other of the node and the user equipment unit serving as a
6 receiver of the segments of data packets;
7 sending an acknowledgment message from the receiver to the transmitter when
8 the receiver detects loss of a segment;
9 characterized by: starting/restarting a poll timer in the transmitter upon
10 occurrence of any of the following:

- 11 (1) a sequence number of a segment to be sent from the transmitter to the
12 receiver is greater than or equal to $V(a) + \text{MaxWin} - \text{MaxP}$, wherein $V(a)$ is a
13 sequence number of a first segment in a transmission window, MaxWin is the
14 size of the transmission window, and MaxP is a predefined integer less than
15 MaxWin;
16 (2) one or more segments are to be retransmitted from the transmitter to the
17 receiver if the poll timer is already running;
18 (3) a segment to be transmitted is the last segment in the buffer.

1 70. The method of claim 69, wherein the integer MaxP is dynamically
2 determined.

1 71. The method of claim 69, wherein the integer MaxP depends on a number of
2 consecutive lost acknowledgment messages allowed.

1 72. The method of claim 71, wherein MaxP is determined by the following
2 expression:


3
$$\text{MaxP} = \text{BW}/\text{PS} * ((1+N) * \text{RTT} + N * \text{POLLT})$$

4 wherein BW is bandwidth; PS is segment size, RTT is round trip time, and POLLT is a
5 poll timer value, and N is set to provide a risk for transmission stall p below a limit.


1 73. The method of claim 69, wherein the integer MaxP depends on a frequency
2 of negative acknowledgment messages.

1 74. The method of claim 69, wherein a time out value for the poll timer is
2 dynamically determined.


1/19

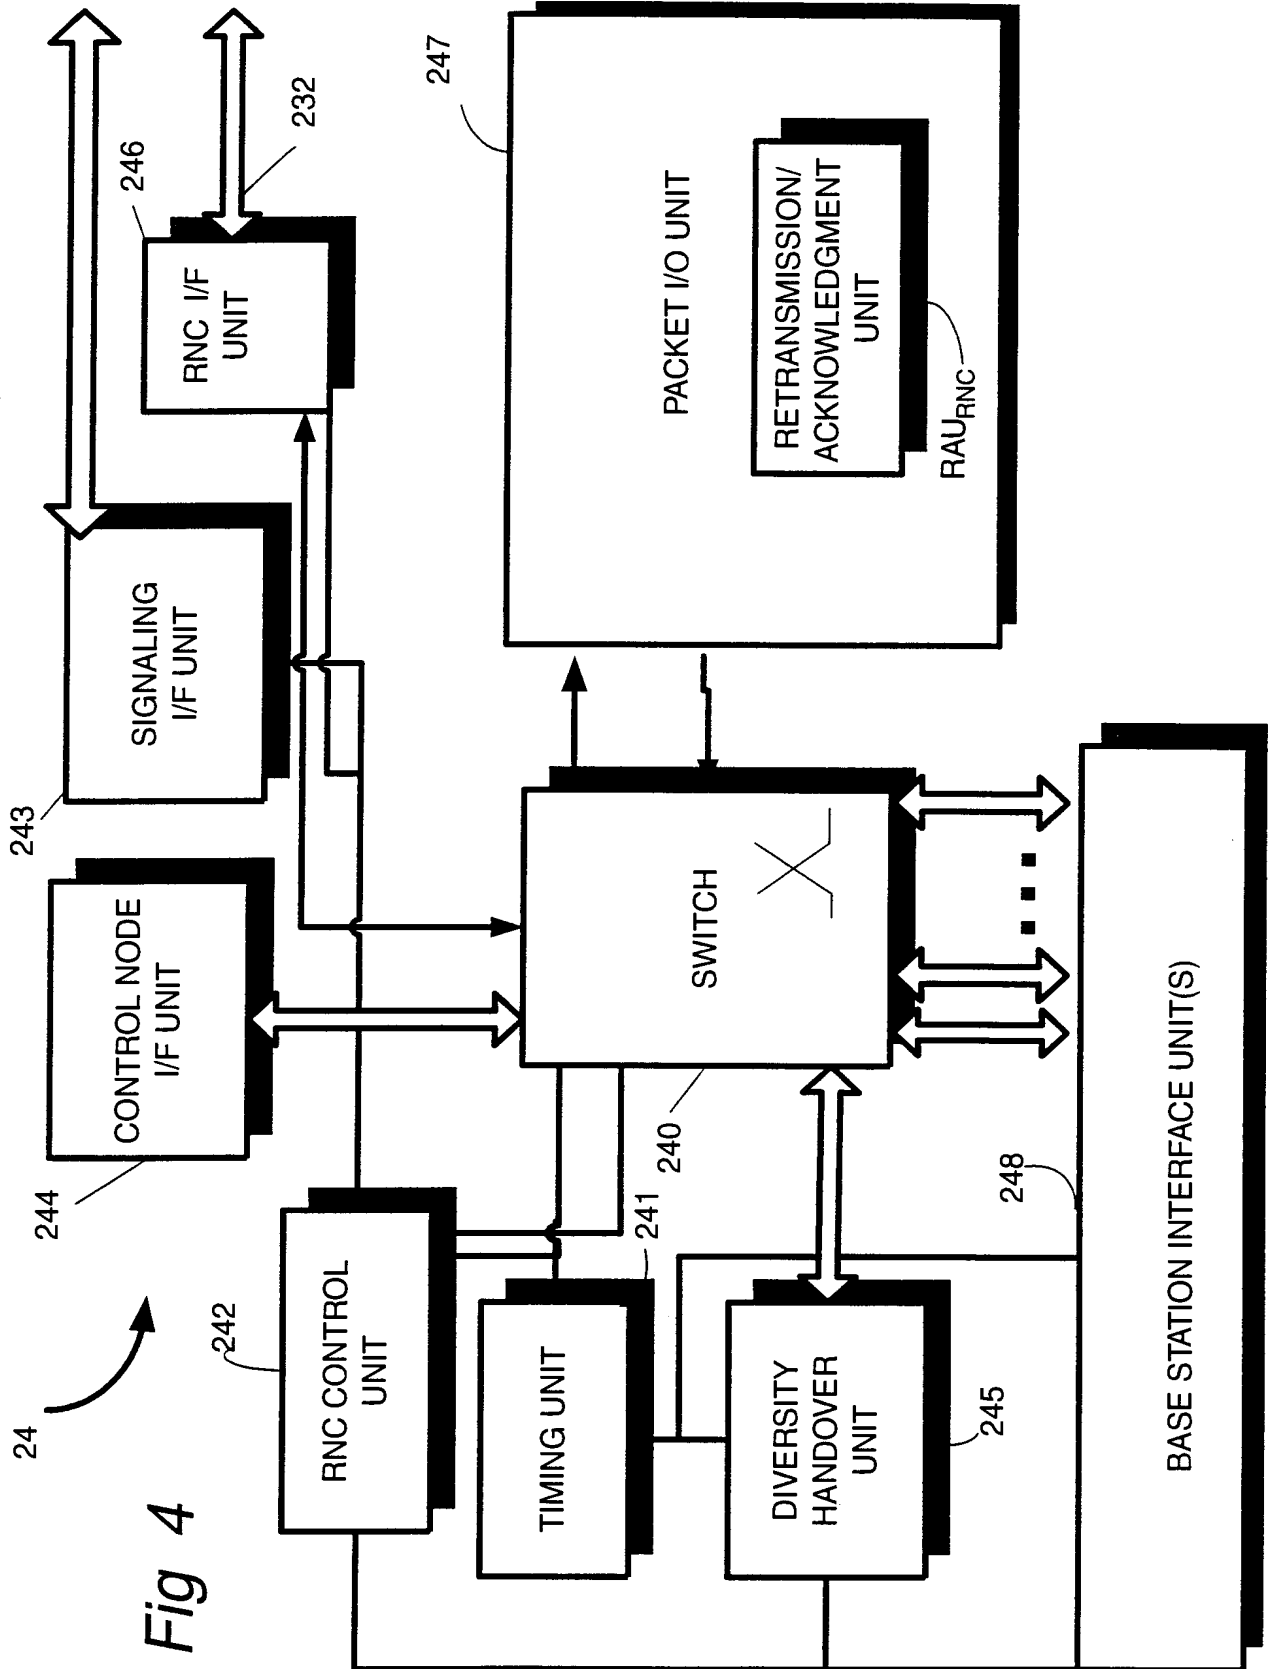


2/19


*Fig 2*

3/19

*Fig 3*



5/19




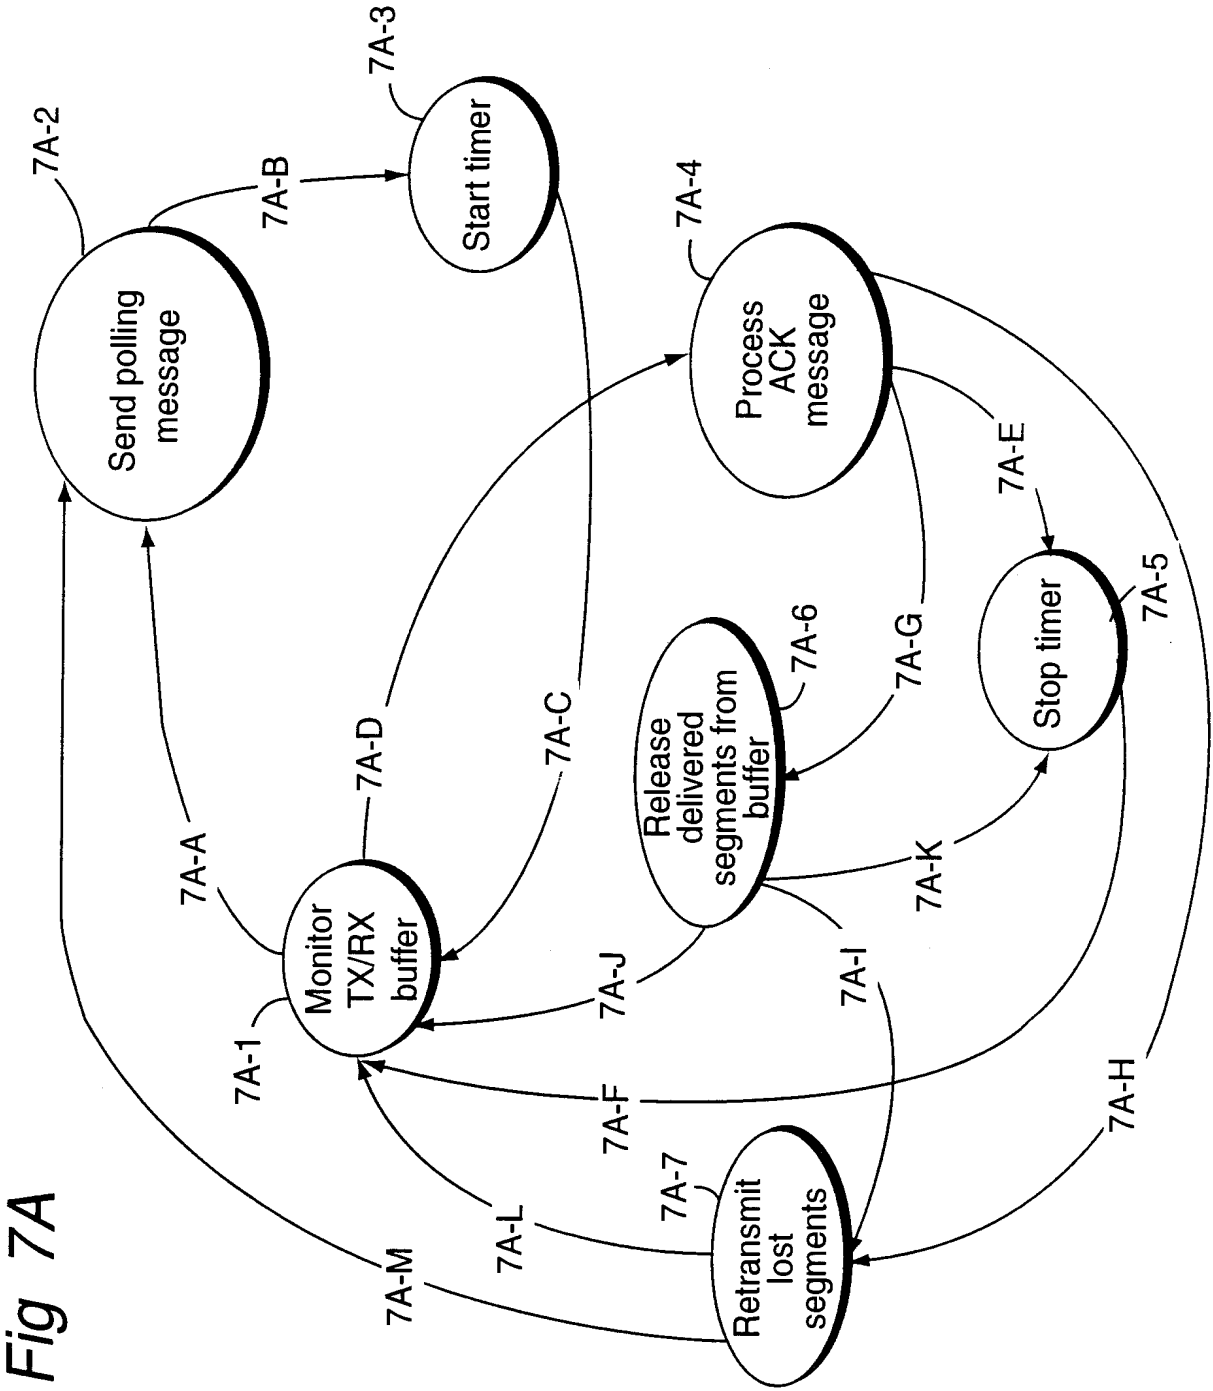
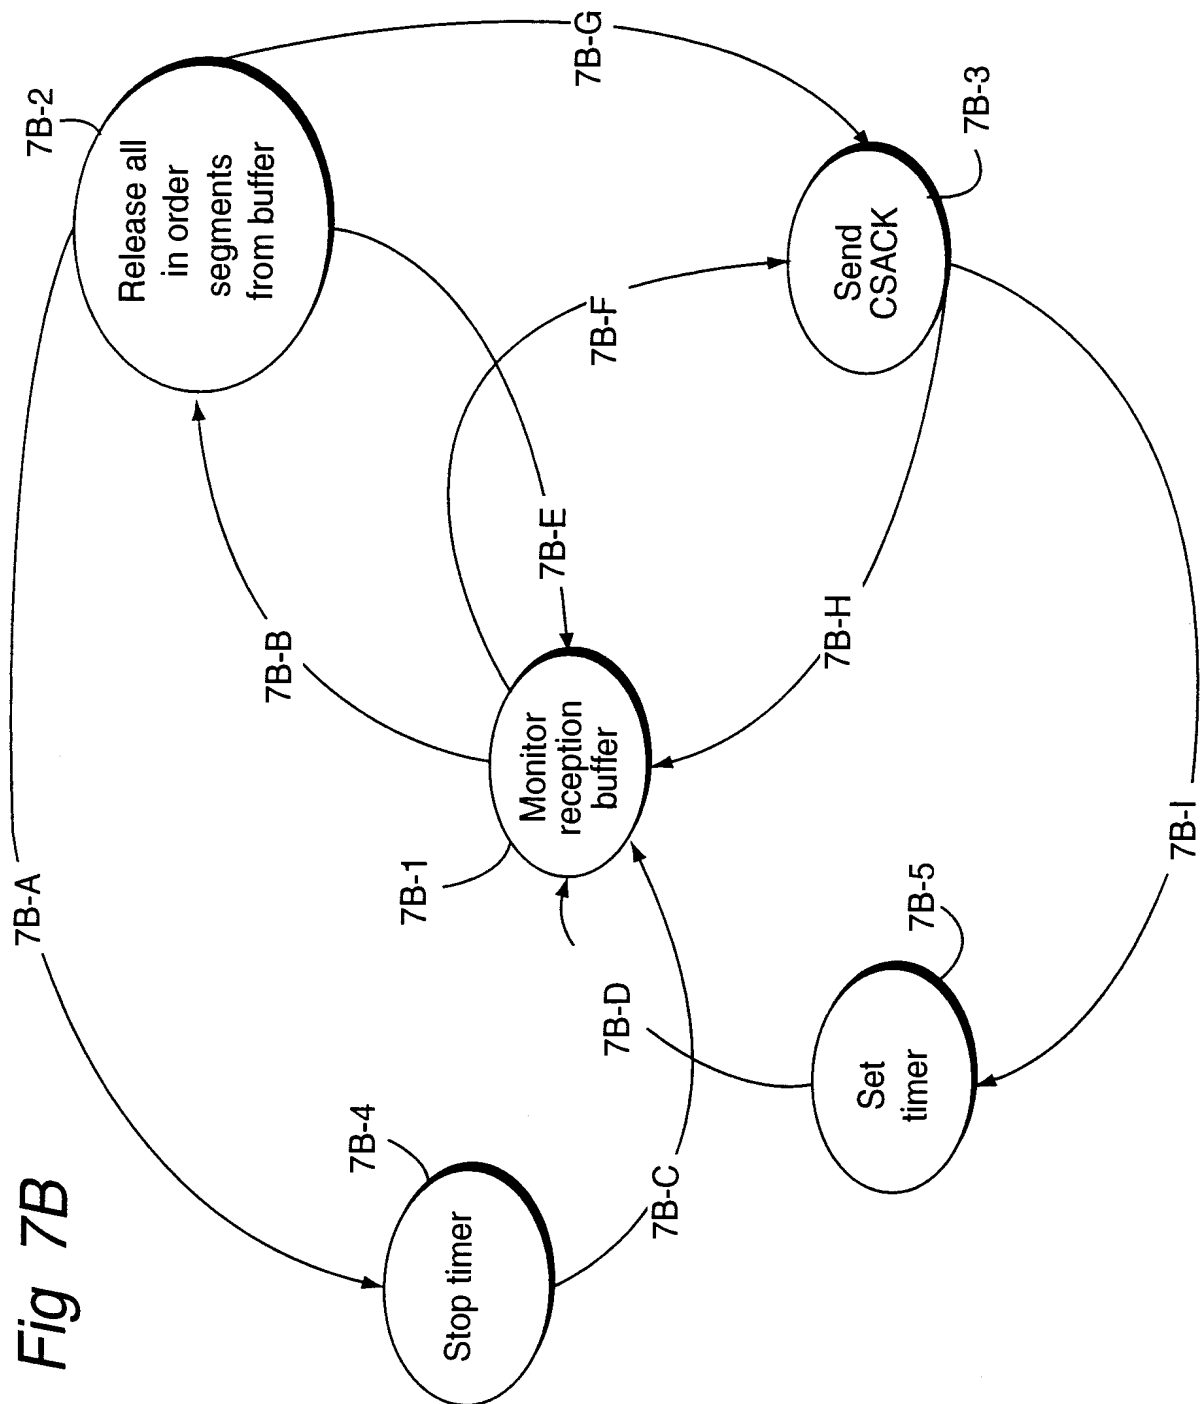


Fig. 6


7/19




8/19



9/19

Fig. 8A

10/19

Fig. 8B

11/19


Fig. 8C

Fig. 9A

C	SSN
SSN	BITMAP
BITMAP	
.	
.	
.	
BITMAP	

OCTET 1

OCTET 2

OCTET 3

OCTET n

Fig. 9B

C	SSN
SSN	BITMAP
BITMAP	
.	
.	
.	
BITMAP	S

OCTET 1

OCTET 2

OCTET 3

OCTET n

Fig. 10B







Fig. 10A




14/19

Fig. 10C

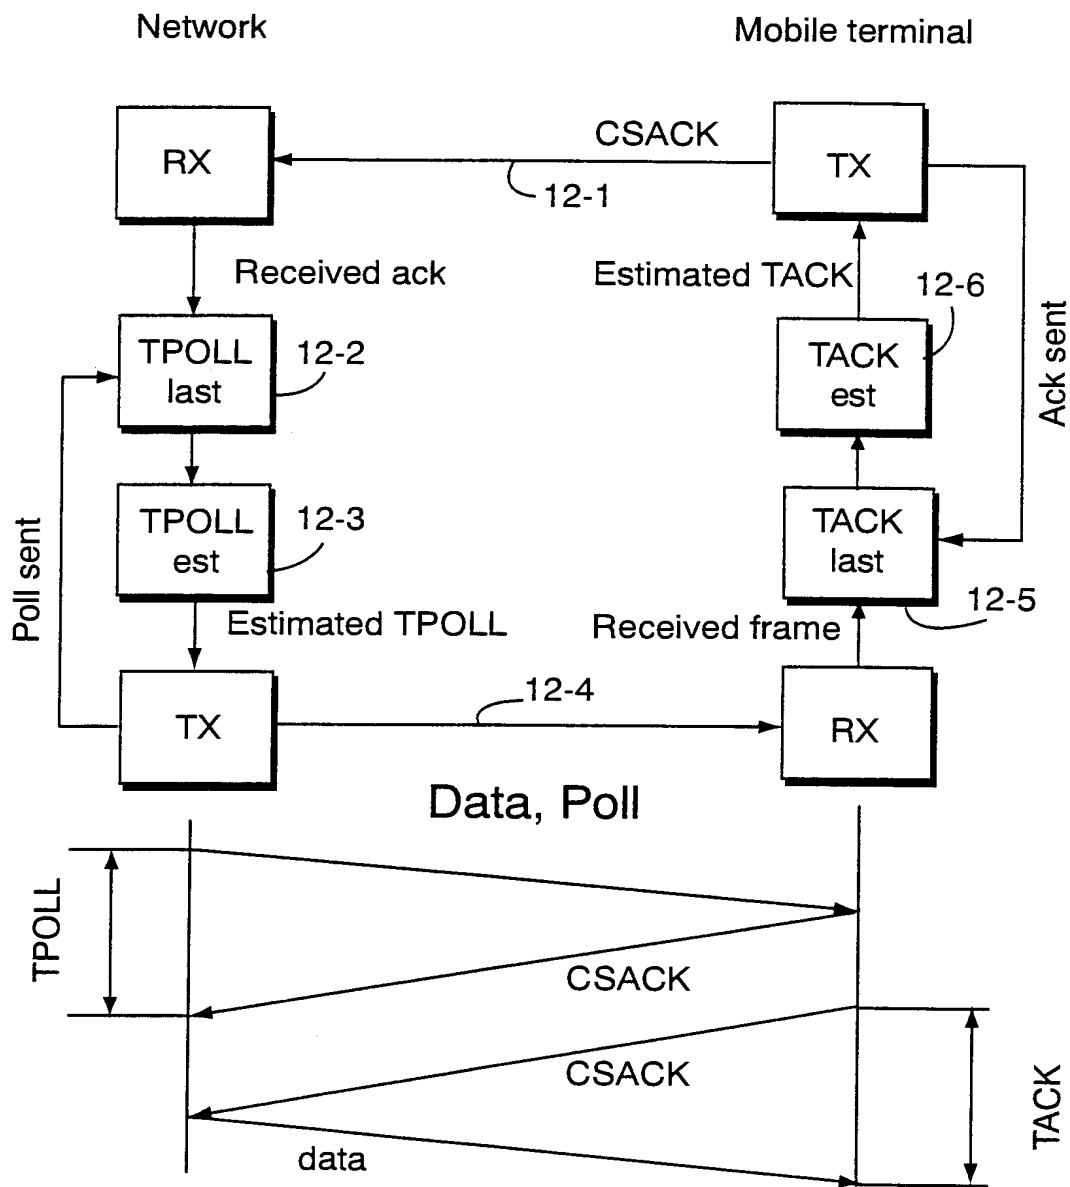
15/19

Fig. 10D



16/19



17/19

Fig. 12A*Fig. 12B*

18/19

Fig. 13A*Fig. 13B*

19/19

BLER	P(stall)	1	2	3	4	5
0.01		0.01	0.0001	0.000001	1E-08	1E-10
0.03	0.03	0.0009	0.000027	8.1E-07	2.43E-08	
0.05	0.05	0.0025	0.000125	6.25E-06	3.13E-07	
0.07	0.07	0.0049	0.000343	2.4E-05	1.68E-06	
0.09	0.09	0.0081	0.000729	6.56E-05	5.9E-06	
0.11	0.11	0.0121	0.001331	0.000146	1.61E-05	
0.13	0.13	0.0169	0.002197	0.000286	3.71E-05	
0.15	0.15	0.0225	0.003375	0.000506	7.59E-05	
0.17	0.17	0.0289	0.004913	0.000835	0.000142	
0.19	0.19	0.0361	0.006859	0.001303	0.000248	
0.21	0.21	0.0441	0.009261	0.001945	0.000408	
0.23	0.23	0.0529	0.012167	0.002798	0.000644	
0.25	0.25	0.0625	0.015625	0.003906	0.000977	
0.27	0.27	0.0729	0.019683	0.005314	0.001435	
0.29	0.29	0.0841	0.024389	0.007073	0.002051	
0.31	0.31	0.0961	0.029791	0.009235	0.002863	
0.33	0.33	0.1089	0.035937	0.011859	0.003914	
0.35	0.35	0.1225	0.042875	0.015006	0.005252	
MaxP		14.336	23.552	32.768	41.984	51.2
RTT		0.05				
BW		32768				
PS		320				
POLLT		0.04				
Limit		0.01				

Fig. 14

INTERNATIONAL SEARCH REPORT

International Application No

PCT/SE 00/02268

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L1/18

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 0 695 053 A (AT & T CORP) 31 January 1996 (1996-01-31) abstract column 1, line 32 -column 2, line 17 column 3, line 54 -column 4, line 27	33, 40
Y	column 5, line 58 -column 6, line 23	34-39, 41-46
X	US 4 975 952 A (MABEY PETER J ET AL) 4 December 1990 (1990-12-04) abstract column 1, line 8 - line 10 column 2, line 51 - line 64 column 5, line 14 -column 6, line 5	47, 55, 63, 69
Y	column 6, line 54 -column 7, line 19	48-54, 56-62, 64-68, 70-74
	--- -/--	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- * & * document member of the same patent family

Date of the actual completion of the international search

21 February 2001

Date of mailing of the international search report

27/02/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Toumpoulidis, T

INTERNATIONAL SEARCH REPORT

International Application No

PCT/SE 00/02268

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	EP 0 658 028 A (AT & T CORP) 14 June 1995 (1995-06-14) abstract column 3, line 24 -column 4, line 56 figures 1-4 ---	34-39, 41-46
Y	ITU-T RECOMMENDATION Q.2110: "B-ISDN ATM Adaptation Layer - Service Specific Connection Oriented Protocol (SSCOP)" July 1994 (1994-07), INTERNATIONAL TELECOMMUNICATION UNION, GENEVA, CH XP002142754 cited in the application page 84, line 1 -page 95, line 37 paragraphs '07.6!', '07.7! figure 5 ---	48-54, 56-62, 64-68, 70-74
A	US 5 440 545 A (HAMILTON JR R LEE ET AL) 8 August 1995 (1995-08-08) abstract column 2, line 67 -column 3, line 11 column 8, line 8 - line 36 column 9, line 16 - line 25 column 10, line 52 -column 11, line 23 column 13, line 17 -column 14, line 27 ---	1-32
A	SIMMONS J M: "Proof of correctness of ATM retransmission scheme" COMPUTER NETWORKS AND ISDN SYSTEMS, NL, NORTH HOLLAND PUBLISHING. AMSTERDAM, vol. 29, no. 2, 1997, pages 181-194, XP004034004 ISSN: 0169-7552 the whole document ---	1-74
P,X	WO 00 33503 A (ERICSSON TELEFON AB L M) 8 June 2000 (2000-06-08) abstract -----	1-4, 17-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/SE 00/02268

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
EP 0695053	A	31-01-1996	US 5570367 A CA 2145782 A,C CN 1123980 A JP 8065355 A	29-10-1996 29-01-1996 05-06-1996 08-03-1996
US 4975952	A	04-12-1990	GB 2180127 A AU 605467 B AU 6210786 A DE 3677430 D DK 415986 A,B, EP 0213682 A ES 2001632 A FI 863520 A,B, JP 2746363 B JP 62059436 A	18-03-1987 17-01-1991 05-03-1987 14-03-1991 05-03-1987 11-03-1987 01-06-1988 05-03-1987 06-05-1998 16-03-1987
EP 0658028	A	14-06-1995	US 5444718 A AU 672536 B AU 7900894 A CA 2132649 A,C CN 1111889 A JP 7202856 A KR 149899 B	22-08-1995 03-10-1996 08-06-1995 31-05-1995 15-11-1995 04-08-1995 15-05-1999
US 5440545	A	08-08-1995	AU 675868 B AU 7321094 A BR 9405540 A CA 2143950 A CN 1113402 A,B EP 0667074 A JP 8502159 T JP 3080991 B KR 197922 B WO 9504421 A	20-02-1997 28-02-1995 08-09-1999 09-02-1995 13-12-1995 16-08-1995 05-03-1996 28-08-2000 15-06-1999 09-02-1995
WO 0033503	A	08-06-2000	AU 2014300 A	19-06-2000